



Cisco Firepower Threat Defense Configuration Guide for Firepower Device Manager, Version 6.2

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

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

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

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

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

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

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

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

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

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

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

© 2015–2018 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1

Getting Started 1

Is This Guide for You?	1
New Features in FDM/FTD 6.2	2
Logging Into the System	4
Logging Into the FDM	4
Logging Into the Command Line Interface (CLI)	5
Changing Your Password	6
Setting User Profile Preferences	6
Creating Local User Accounts for the FTD CLI	7
Setting Up the System	9
Connect the Interfaces	9
Cabling for ASA 5506-X, 5506W-X, and 5506H-X	10
Cabling for ASA 5508-X and 5516-X	11
Cabling for ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X	12
(Optional) Change Management Network Settings at the CLI	12
Complete the Initial Configuration Using the Setup Wizard	14
What to Do if the Outside Subnet Conflicts with the Inside Subnet (Setup Wizard Hangs at Step 1)	16
Configure the Wireless Access Point (ASA 5506W-X)	18
Default Configuration Prior to Initial Setup	21
Configuration After Initial Setup	22
Configuration Basics	25
Configuring the Device	25
Configuring Security Policies	26
Deploying Your Changes	27
Configuration Changes that Restart Inspection Engines	28

Viewing Interface and Management Status 28

Viewing System Task Status 29

CHAPTER 2**Best Practices: Use Cases for FTD 31**

How to Configure the Device in FDM 31

How to Gain Insight Into Your Network Traffic 36

How to Block Threats 42

How to Block Malware 46

How to Implement an Acceptable Use Policy (URL Filtering) 49

How to Control Application Usage 53

How to Add a Subnet 57

More Examples 62

CHAPTER 3**Licensing the System 63**

Smart Licensing for the Firewall System 63

Cisco Smart Software Manager 63

Periodic Communication with the License Authority 64

Smart License Types 64

Impact of Expired or Disabled Optional Licenses 65

Managing Smart Licenses 65

Registering the Device 66

Enabling or Disabling Optional Licenses 67

Synchronizing with the Cisco Smart Software Manager 67

Unregistering the Device 68

CHAPTER 4**Monitoring the Device 69**

Enable Logging to Obtain Traffic Statistics 69

Event Types 69

Configurable Connection Logging 70

Automatic Connection Logging 70

Tips for Connection Logging 70

Sending Events to an External Syslog Server 71

Monitoring Traffic and System Dashboards 71

Monitoring Additional Statistics Using the Command Line 73

Viewing Events	74
Configuring Custom Views	75
Filtering Events	76
Event Field Descriptions	77

CHAPTER 5**Objects 85**

Object Types	85
Managing Objects	86
Configuring Network Objects and Groups	87
Configuring Port Objects and Groups	88
Configuring Security Zones	89
Configuring Application Filter Objects	90
Configuring URL Objects and Groups	92
Configuring Geolocation Objects	94
Configuring Syslog Servers	94

PART I**The Basics 97****CHAPTER 6****Interfaces 99**

About FTD Interfaces	99
Interface Modes	99
Management/Diagnostic Interface	100
Recommendations for Configuring a Separate Management Network	100
Limitations for Management/Diagnostic Interface Configuration for a Separate Management Network	101
Security Zones	101
IPv6 Addressing	101
Auto-MDI/MDIX Feature	102
Guidelines and Limitations for Interfaces	102
Limitations for Interface Configuration	102
Maximum Number of VLAN Subinterfaces by Device Model	103
Configure a Physical Interface	104
Configure Bridge Groups	107
Configure VLAN Subinterfaces and 802.1Q Trunking	110

- Configure Advanced Interface Options 113
 - About the MTU 113
 - Path MTU Discovery 114
 - MTU and Fragmentation 114
 - MTU and Jumbo Frames 114
 - Configure Advanced Options 114
- Monitoring Interfaces 116
- Examples for Interfaces 116

CHAPTER 7

Routing 117

- Routing Overview 117
 - The Routing Table and Route Selection 117
 - How Forwarding Decisions Are Made 117
 - Routing Table for Management Traffic 118
 - Equal-Cost Multi-Path (ECMP) Routing 119
- Static Routes 119
 - About Static and Default Routes 119
 - Default Route 119
 - Static Routes 120
 - Guidelines for Static Routing 120
 - Configuring Static Routes 120
- Monitoring Routing 121

PART II

Security Policies 123

CHAPTER 8

Identity Policies 125

- Identity Policy Overview 125
 - Establishing User Identity through Active Authentication 125
 - Limitations on Number of Users 126
 - Supported Directory Servers 126
 - Determining the Directory Base DN 126
 - Dealing with Unknown Users 127
- Configuring Identity Policies 128
 - Configure Directory Server 129

Configure the Active Authentication Captive Portal	130
Configure Identity Rules	131
Enabling Transparent User Authentication	134
Requirements for Transparent Authentication	135
Configuring Internet Explorer for Transparent Authentication	135
Configuring Firefox for Transparent Authentication	136
Monitoring Identity Policies	137
Examples for Identity Policies	137

CHAPTER 9
Access Control 139

Access Control Overview	139
Access Control Rules and the Default Action	139
Application Filtering	140
Application Control for Encrypted and Decrypted Traffic	140
Best Practices for Application Filtering	140
URL Filtering	141
Filtering URLs by Category and Reputation	141
Looking Up the Category and Reputation for a URL	142
Manual URL Filtering	142
Filtering HTTPS Traffic	143
Comparing URL and Application Filtering	144
Best Practices for Effective URL Filtering	144
What the User Sees When You Block Web Sites	144
Intrusion, File, and Malware Inspection	145
Best Practices for Access Control Rule Order	146
NAT and Access Rules	146
How Other Security Policies Impact Access Control	146
License Requirements for Access Control	146
Guidelines and Limitations for Access Control Policies	147
Configuring the Access Control Policy	148
Configuring the Default Action	148
Configuring Access Control Rules	149
Source/Destination Criteria	150
Application Criteria	152

URL Criteria	153
User Criteria	154
Intrusion Policy Settings	155
File Policy Settings	156
Logging Settings	157
Monitoring Access Control Policies	158
Monitoring Access Control Statistics in the Dashboards	158
Monitoring Access Control Policies in the CLI	158
Examples for Access Control	159

CHAPTER 10
Network Address Translation (NAT) 161

Why Use NAT?	161
NAT Basics	162
NAT Terminology	162
NAT Types	162
NAT in Routed Mode	163
Auto NAT and Manual NAT	163
Auto NAT	164
Manual NAT	164
Comparing Auto NAT and Manual NAT	164
NAT Rule Order	165
NAT Interfaces	167
Configuring Routing for NAT	167
Addresses on the Same Network as the Mapped Interface	167
Addresses on a Unique Network	168
The Same Address as the Real Address (Identity NAT)	168
Guidelines for NAT	168
Interface Guidelines	168
IPv6 NAT Guidelines	168
IPv6 NAT Best Practices	169
NAT Support for Inspected Protocols	169
Additional Guidelines for NAT	171
Configure NAT	173
Dynamic NAT	173

About Dynamic NAT	173
Dynamic NAT Disadvantages and Advantages	174
Configure Dynamic Auto NAT	175
Configure Dynamic Manual NAT	176
Dynamic PAT	178
About Dynamic PAT	178
Dynamic PAT Disadvantages and Advantages	179
Configure Dynamic Auto PAT	179
Configure Dynamic Manual PAT	180
Static NAT	182
About Static NAT	183
Configure Static Auto NAT	186
Configure Static Manual NAT	188
Identity NAT	191
Configure Identity Auto NAT	191
Configure Identity Manual NAT	192
NAT Rule Properties for FTD	194
Packet Translation Properties for Auto NAT	195
Packet Translation Properties for Manual NAT	196
Advanced NAT Properties	197
Translating IPv6 Networks	198
NAT64/46: Translating IPv6 Addresses to IPv4	199
NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet	199
NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet and DNS Translation	201
NAT66: Translating IPv6 Addresses to Different IPv6 Addresses	206
NAT66 Example, Static Translation between Networks	206
NAT66 Example, Simple IPv6 Interface PAT	209
Monitoring NAT	212
Examples for NAT	213
Providing Access to an Inside Web Server (Static Auto NAT)	213
Single Address for FTP, HTTP, and SMTP (Static Auto NAT-with-Port-Translation)	215
Different Translation Depending on the Destination (Dynamic Manual PAT)	221
Different Translation Depending on the Destination Address and Port (Dynamic Manual PAT)	227
Rewriting DNS Queries and Responses Using NAT	232

DNS 64 Reply Modification	233
DNS Reply Modification, DNS Server on Outside	239
DNS Reply Modification, DNS Server on Host Network	242

PART III
Virtual Private Networks (VPN) 247

CHAPTER 11
Site-to-Site VPN 249

VPN Basics	249
Internet Key Exchange (IKE)	250
How Secure Should a VPN Connection Be?	250
Deciding Which Encryption Algorithm to Use	251
Deciding Which Hash Algorithms to Use	252
Deciding Which Diffie-Hellman Modulus Group to Use	252
VPN Topologies	253
Managing Site-to-Site VPNs	253
Configuring a Site-to-Site VPN Connection	254
Configuring the Global IKE Policy	256
Configuring IKEv1 Policies	258
Configuring IKEv2 Policies	259
Configuring IPsec Proposals	261
Configuring IPsec Proposals for IKEv1	261
Configuring IPsec Proposals for IKEv2	262
Verifying Site-to-Site VPN Connections	263
Monitoring Site-to-Site VPN	266
Examples for Site-to-Site VPN	266
Exempting Site-to-Site VPN Traffic from NAT	266
How to Provide Internet Access on the Outside Interface for External Site-to-Site VPN Users (Hair Pinning)	272

PART IV
System Administration 279

CHAPTER 12
System Settings 281

Configuring the Management Access List	281
Configuring Diagnostic Logging	283

Severity Levels	283
Configuring the DHCP Server	284
Configuring DNS	286
Configuring the Management Interface	286
Configuring the Device Hostname	288
Configuring Network Time Protocol (NTP)	288
Configuring URL Filtering Preferences	289
Configuring Cloud Management (Cisco Defense Orchestrator)	289

CHAPTER 13
System Management 291

Installing Software Updates	291
Updating System Databases	291
Overview of System Database Updates	291
Updating System Databases	292
Upgrading FTD Software	293
Reimaging the Device	295
Backing Up and Restoring the System	295
Backing Up the System Immediately	296
Backing Up the System at a Scheduled Time	297
Setting Up a Recurring Backup Schedule	297
Restoring a Backup	298
Managing Backup Files	299
Rebooting the System	299
Troubleshooting the System	300
Pinging Addresses to Test Connectivity	300
Tracing Routes to Hosts	302
Troubleshooting NTP	303
Troubleshooting DNS for the Management Interface	305
Analyzing CPU and Memory Usage	307
Viewing Logs	308
Creating a Troubleshooting File	309
Uncommon Management Tasks	310
Switching Between Local and Remote Management	310
Changing the Firewall Mode	312

[Resetting the Configuration](#) 315



CHAPTER 1

Getting Started

The following topics explain how to get started configuring the Firepower Threat Defense (FTD).

- [Is This Guide for You?, on page 1](#)
- [New Features in FDM/FTD 6.2, on page 2](#)
- [Logging Into the System, on page 4](#)
- [Setting Up the System, on page 9](#)
- [Configuration Basics, on page 25](#)

Is This Guide for You?

This guide explains how to configure FTD using the Firepower Device Manager (FDM) web-based configuration interface included on the FTD devices.

The FDM lets you configure the basic features of the software that are most commonly used for small or mid-size networks. It is especially designed for networks that include a single device or just a few, where you do not want to use a high-powered multiple-device manager to control a large network containing many FTD devices.

If you are managing large numbers of devices, or if you want to use the more complex features and configurations that FTD allows, use the Firepower Management Center (FMC) to configure your devices instead of the integrated FDM.

You can use the FDM on the following devices.

Table 1: FDM Supported Models

Device Model	Minimum FTD Software Version
ASA 5508-X, 5516-X	6.1
ASA 5525-X, 5545-X, 5555-X	6.1
ASA 5506-X, 5506H-X, 5506W-X, 5512-X	6.1
ASA 5515-X	6.1

New Features in FDM/FTD 6.2

Released: January 23, 2017

The following table lists the new features available in FTD 6.2 when configured using FDM.

Feature	Description
Cisco Defense Orchestrator (CDO) cloud management.	You can manage the device using the Cisco Defense Orchestrator cloud-based portal. Select Device > System Settings > Cloud Management . For more information on Cisco Defense Orchestrator, see http://www.cisco.com/go/cdo .
Drag and drop for access rules.	You can drag and drop access rules to move them in the rules table.
Upgrade FTD software through FDM.	You can install software upgrades through FDM. Select Device > Updates .
Default configuration changes.	For new or reimaged devices, the default configuration includes significant changes, including: <ul style="list-style-type: none"> • (ASA 5506-X, 5506W-X, 5506H-X.) Except for the first data interface, and the Wi-Fi interface on an ASA 5506W-X, all other data interfaces on these device models are structured into the “inside” bridge group and enabled. There is a DHCP server on the inside bridge group. You can plug endpoints or switches into any bridged interface and endpoints get addresses on the 192.168.1.0/24 network. • The inside interface IP address is now 192.168.1.1, and a DHCP server is defined on the interface with the address pool 192.168.1.5-192.168.1.254. • HTTPS access is enabled on the inside interface, so you can open FDM through the inside interface at the default address, 192.168.1.1. For the ASA 5506-X models, you can do this through any inside bridge group member interface. • The management port hosts a DHCP server for the 192.168.45.0/24 network. You can plug a workstation directly into the management port, get an IP address, and open FDM to configure the device. • The OpenDNS public DNS servers are now the default DNS servers for the management interface. Previously, there were no default DNS servers. You can configure different DNS servers during device setup. • The default gateway for the management IP address is to use the data interfaces to route to the Internet. Thus, you do not need to wire the Management physical interface to a network.

Feature	Description
Management interface and access changes.	<p>Several changes to how the management address, and access to FDM, works:</p> <ul style="list-style-type: none"> • You can now open data interfaces to HTTPS (for FDM) and SSH (for CLI) connections. You do not need a separate management network, or to connect the Management/Diagnostic physical port to the inside network, to manage the device. Select Device > System Settings > Management Access List. • The system can obtain system database updates through the gateway for the outside interface. You do not need to have an explicit route from the management interface or network to the Internet. The default is to use internal routes through the data interfaces. However, you can set a specific gateway if you prefer to use a separate management network. Select Device > System Settings > Management Interface. • You can use FDM to configure the management interface to obtain its IP address through DHCP. Select Device > System Settings > Management Interface. • You can configure a DHCP server on the management address if you configure a static address. Select Device > System Settings > Management Interface.
Miscellaneous user interface changes.	<p>The following are notable changes to the FDM user interface.</p> <ul style="list-style-type: none"> • Device main menu item. In previous releases, this menu item was the host name of your device. Also, the page opened is called Device Summary instead of Device Dashboard. • You cannot select an alternative outside interface during initial device setup. The first data interface is the default outside interface. • Device > System Settings > Cloud Preferences is now called Device > System Settings > URL Filtering Preferences. • The System Settings > DHCP Server page is now organized on two tabs, with the table of DHCP servers separated from the global parameters.
Site-to-site VPN connections.	<p>You can configure site-to-site virtual private network (VPN) connections using preshared keys. You can configure IKEv1 and IKEv2 connections.</p>

Feature	Description
Integrated Routing and Bridging support.	<p>Integrated Routing and Bridging provides the ability to route between a bridge group and a routed interface. A bridge group is a group of interfaces that the FTD device bridges instead of routes. The FTD device is not a true bridge in that the FTD device continues to act as a firewall: access control between interfaces is controlled, and all of the usual firewall checks are in place.</p> <p>This feature lets you configure bridge groups and to route between bridge groups and between a bridge group and a routed interface. The bridge group participates in routing by using a Bridge Virtual Interface (BVI) to act as a gateway for the bridge group. Integrated Routing and Bridging provides an alternative to using an external Layer 2 switch if you have extra interfaces on the FTD device to assign to the bridge group. The BVI can be a named interface and can participate separately from member interfaces in some features, such as DHCP server, where you configure other features on bridge group member interfaces, such as NAT and access control rules.</p> <p>Select Device > Interfaces to configure a bridge group.</p>

Logging Into the System

There are two interfaces to the FTD device:

FDM Web Interface

The FDM runs in your web browser. You use this interface to configure, manage, and monitor the system.

Command Line Interface (CLI, Console)

Use the CLI for troubleshooting. You can also use it for initial setup instead of the FDM.

The following topics explain how to log into these interfaces and manage your user account.

Logging Into the FDM

Use the FDM to configure, manage, and monitor the system. The features that you can configure through the browser are not configurable through the command-line interface (CLI); you must use the web interface to implement your security policies.

Use a current version of the following browsers: Firefox, Chrome, Safari, Edge, or Internet Explorer.



Note If you type in the wrong password and fail to log in on 3 consecutive attempts, your account is locked for 5 minutes. You must wait before trying to log in again.

Before you begin

You can log into the FDM using the **admin** username only. You cannot create additional users for the FDM access.

There can be up to 5 active logins at one time. This includes users logged into the device manager and active API sessions, which are represented by non-expired API tokens. If you exceed this limit, the oldest session, either the device manager login or API token, is expired to allow the new session. These limits do not apply to SSH sessions.

Procedure

Step 1 Using a browser, open the home page of the system, for example, <https://ftd.example.com>.

You can use any of the following addresses. You can use the IPv4 or IPv6 address or the DNS name, if you have configured one.

- The management address. By default (on most platforms), this is 192.168.45.45 on the Management interface.
- The address of a data interface that you have opened for HTTPS access. By default, the “inside” interface allows HTTPS access, so you can connect to the default inside address 192.168.1.1. On device models where the inside interface is a bridge group, you can connect to this address through any bridge group member interface. See [Default Configuration Prior to Initial Setup, on page 21](#) for details about your model's inside IP address.

Tip If your browser is not configured to recognize the server certificate, you will see a warning about an untrusted certificate. Accept the certificate as an exception, or in your trusted root certificate store.

Step 2 Enter the **admin** username and password, then click **Login**.

The default admin password is Admin123.

Your session will expire after 30 minutes of inactivity, and you will be prompted to log in again. You can log out by selecting **Log Out** from the user icon drop-down menu in the upper right of the page.



Logging Into the Command Line Interface (CLI)

Use the command-line interface (CLI) to set up the system and do basic system troubleshooting. You cannot configure policies through a CLI session.

To log into the CLI, do one of the following:

- Use the console cable included with the device to connect your PC to the console using a terminal emulator set for 9600 baud, 8 data bits, no parity, 1 stop bit, no flow control. See the hardware guide for your device for more information about the console cable.
- Use an SSH client to make a connection to the management IP address. You can also connect to the address on a data interface if you open the interface for SSH connections (see [Configuring the Management Access List, on page 281](#)). SSH access to data interfaces is disabled by default. Log in using the **admin** username or another CLI user account. The default admin password is Admin123.

Tips

- After logging in, for information on the commands available in the CLI, enter **help** or **?**. For usage information, see [Cisco Firepower Threat Defense Command Reference](http://www.cisco.com/c/en/us/td/docs/security/firepower/command_ref/b_Command_Reference_for_Firepower_Threat_Defense.html) at http://www.cisco.com/c/en/us/td/docs/security/firepower/command_ref/b_Command_Reference_for_Firepower_Threat_Defense.html.
- You can create local user accounts that can log into the CLI using the **configure user add** command. However, these users can log into the CLI only. They cannot log into the FDM web interface.

Changing Your Password

You should periodically change your password. The following procedure explains how to change the password while logged into FDM.



Note If you are logged into the CLI, you can change your password using the **configure password** command. You can change the password for a different CLI user with the **configure user password *username*** command.

Procedure

Step 1 Select **Profile** from the user icon drop-down list in the upper right of the menu.



Step 2 Click the **Password** tab.

Step 3 Enter your current password.

Step 4 Enter your new password and then confirm it.

Step 5 Click **Change**.

Setting User Profile Preferences

You can set preferences for the user interface and change your password.

Procedure

Step 1 Select **Profile** from the user icon drop-down list in the upper right of the menu.



Step 2 On the **Profile** tab, configure the following and click **Save**.

- **Time Zone for Scheduling Tasks**—Select the time zone you want to use for scheduling tasks such as backups and updates. The browser time zone is used for dashboards and events, if you set a different zone.
- **Color Theme**—Select the color theme you want to use in the user interface.

Step 3 On the **Password** tab, you can enter a new password and click **Change**.

Creating Local User Accounts for the FTD CLI

You can create users for CLI access on FTD devices. These accounts do not allow access to the management application, but to the CLI only. The CLI is useful for troubleshooting and monitoring purposes.

You cannot create local user accounts on more than one device at a time. Each device has its own set of unique local user CLI accounts.

Procedure

Step 1 Log into the device CLI using an account with config privileges.

The admin user account has the required privileges, but any account with config privileges will work. You can use an SSH session or the Console port.

For certain device models, the Console port puts you into the FXOS CLI. Use the **connect ftd** command to get to the FTD CLI.

Step 2 Create the user account.

configure user add *username* {**basic** | **config**}

You can define the user with the following privilege levels:

- **config**—Gives the user configuration access. This gives the user full administrator rights to all commands.
- **basic**—Gives the user basic access. This does not allow the user to enter configuration commands.

Example:

The following example adds a user account named joecool with config access rights. The password is not shown as you type it.

```
> configure user add joecool config
Enter new password for user joecool: newpassword
Confirm new password for user joecool: newpassword
> show user
Login          UID   Auth Access  Enabled Reset   Exp Warn  Str Lock Max
admin          1000 Local Config Enabled  No   Never  N/A  Dis  No  N/A
joecool        1001 Local Config Enabled  No   Never  N/A  Dis  No   5
```

Note Tell users they can change their passwords using the **configure password** command.

Step 3 (Optional.) Adjust the characteristics of the account to meet your security requirements.

You can use the following commands to change the default account behavior.

- **configure user aging** *username max_days warn_days*

Sets an expiration date for the user's password. Specify the maximum number of days for the password to be valid followed by the number of days before expiration the user will be warned about the upcoming expiration. Both values are 1 to 9999, but the warning days must be less than the maximum days. When you create the account, there is no expiration date for the password.

- **configure user forcereset** *username*

Forces the user to change the password on the next login.

- **configure user maxfailedlogins** *username number*

Sets the maximum number of consecutive failed logins you will allow before locking the account, from 1 to 9999. Use the **configure user unlock** command to unlock accounts. The default for new accounts is 5 consecutive failed logins.

- **configure user minpasswdlen** *username number*

Sets a minimum password length, which can be from 1 to 127.

- **configure user strengthcheck** *username {enable | disable}*

Enables or disables password strength checking, which requires a user to meet specific password criteria when changing their password. When a user's password expires or if the **configure user forcereset** command is used, this requirement is automatically enabled the next time the user logs in.

Step 4 Manage user accounts as necessary.

Users can get locked out of their accounts, or you might need to remove accounts or fix other issues. Use the following commands to manage the user accounts on the system.

- **configure user access** *username {basic | config}*

Changes the privileges for a user account.

- **configure user delete** *username*

Deletes the specified account.

- **configure user disable** *username*

Disables the specified account without deleting it. The user cannot log in until you enable the account.

- **configure user enable** *username*

Enables the specified account.

- **configure user password** *username*

Changes the password for the specified user. Users should normally change their own password using the **configure password** command.

- **configure user unlock** *username*

Unlocks a user account that was locked due to exceeding the maximum number of consecutive failed login attempts.

Setting Up the System

You must complete an initial configuration to make the system function correctly in your network. Successful deployment includes attaching cables correctly and configuring the addresses needed to insert the device into your network and connect it to the Internet or other upstream router. The following procedure explains the process.

Before you begin

Before you start the initial setup, the device includes some default settings. For details, see [Default Configuration Prior to Initial Setup, on page 21](#).

Procedure

- Step 1** [Connect the Interfaces, on page 9](#)
 - Step 2** [Complete the Initial Configuration Using the Setup Wizard, on page 14](#)
For details about the resulting configuration, see [Configuration After Initial Setup, on page 22](#).
 - Step 3** [Configure the Wireless Access Point \(ASA 5506W-X\), on page 18](#)
-

Connect the Interfaces

The default configuration assumes that certain interfaces are used for the inside and outside networks. Initial configuration will be easier to complete if you connect network cables to the interfaces based on these expectations.

The default configuration is designed to let you attach your management computer to the inside interface. Alternatively, you can also directly attach your workstation to the Management port. The interfaces are on different networks, so do not try to connect any of the inside interfaces and the Management port to the same network.

Do not connect any of the inside interfaces or the Management interface to a network that has an active DHCP server. This will conflict with the DHCP servers already running on the inside interface and Management interface. If you want to use a different DHCP server for the network, disable the unwanted DHCP server after initial setup.

The following topics show how to cable the system for this topology when using the inside interfaces to configure the device.

Cabling for ASA 5506-X, 5506W-X, and 5506H-X

Figure 1: ASA 5506W-X (with Wi-Fi), 5506-X (without Wi-Fi)





Figure 2: ASA 5506H-X



- Attach GigabitEthernet 1/1 to the ISP/WAN modem or other outside device. By default, the IP address is obtained using DHCP, but you can set a static address during initial configuration.
- Attach GigabitEthernet 1/2 (or another of the inside bridge group member ports) to your workstation, the one you will use to configure the device. Configure the workstation to obtain an IP address using DHCP. The workstation gets an address on the 192.168.1.0/24 network.




Note You have a couple of other options for connecting the management workstation. You can also directly connect it to the Management port. The workstation gets an address through DHCP on the 192.168.45.0/24 network. Another option is to leave your workstation attached to a switch, and attach that switch to one of the inside ports such as GigabitEthernet1/2. However, you must ensure that no other device on the switch's network is running a DHCP server, because it will conflict with the one running on the inside bridge group, 192.168.1.1.

- Optionally, attach other endpoints or switches to the other ports in the inside bridge group. You might want to wait until you complete the initial device setup before adding endpoints. If you add switches, ensure that there are no other DHCP servers running on those networks, as this conflicts with the DHCP server running on the inside bridge group.

Cabling for ASA 5508-X and 5516-X

Figure 3: Cabling the ASA 5508-X or 5516-X



- Connect your management computer to either of the following interfaces:
 - GigabitEthernet 1/2—Connect your management computer directly to GigabitEthernet 1/2 for initial configuration, or connect GigabitEthernet 1/2 to your inside network. GigabitEthernet 1/2 has a default IP address (192.168.1.1) and also runs a DHCP server to provide IP addresses to clients (including the management computer), so make sure these settings do not conflict with any existing inside network settings
 - Management 1/1—Connect your management computer directly to Management 1/1 for initial configuration, or connect Management 1/1 to your management network. Management 1/1 has a default IP address (192.168.45.45) and also runs a DHCP server to provide IP addresses to clients (including the management computer), so make sure these settings do not conflict with any existing inside network settings.

If you need to change the Management 1/1 IP address from the default, you must also cable your management PC to the console port. See [\(Optional\) Change Management Network Settings at the CLI, on page 12](#).

You can later configure the FDM management access from other interfaces.


- Connect the outside network to the GigabitEthernet1/1 interface.

By default, the IP address is obtained using IPv4 DHCP, but you can set a static address during initial configuration.

- Connect other networks to the remaining interfaces.

Cabling for ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X

Figure 4: Cabling the ASA 5500-X



- Connect your management computer to either of the following interfaces:
 - **GigabitEthernet 0/1**—Connect your management computer directly to GigabitEthernet 0/1 for initial configuration, or connect GigabitEthernet 0/1 to your inside network. GigabitEthernet 0/1 has a default IP address (192.168.1.1) and also runs a DHCP server to provide IP addresses to clients (including the management computer), so make sure these settings do not conflict with any existing inside network settings.
 - **Management 0/0**—Connect your management computer directly to Management 0/0 for initial configuration, or connect Management 0/0 to your management network. Management 0/0 has a default IP address (192.168.45.45) and also runs a DHCP server to provide IP addresses to clients (including the management computer), so make sure these settings do not conflict with any existing inside network settings.

If you need to change the Management 0/0 IP address from the default, you must also cable your management computer to the console port. See [\(Optional\) Change Management Network Settings at the CLI](#), on page 12.

You can later configure the FDM management access from other interfaces.

- Connect the outside network to the GigabitEthernet 0/0 interface.
By default, the IP address is obtained using DHCP, but you can set a static address during initial configuration.
- Connect other networks to the remaining interfaces.

(Optional) Change Management Network Settings at the CLI

If you cannot use the default management IP address, then you can connect to the console port and perform initial setup at the CLI, including setting the Management IP address, gateway, and other basic networking settings. You can only configure the Management interface settings; you cannot configure inside or outside interfaces, which you can later configure in the GUI.



Note You cannot repeat the CLI setup script unless you clear the configuration; for example, by reimaging. However, all of these settings can be changed later at the CLI using **configure network** commands. See [Cisco Secure Firewall Threat Defense Command Reference](#).

Procedure

- Step 1** Connect to the FTD console port. See [Logging Into the Command Line Interface \(CLI\)](#), on page 5 for more information.
- Step 2** Log in with the username **admin**.
- The default admin password is Admin123.
- Step 3** The first time you log into the FTD, you are prompted to accept the End User License Agreement (EULA). You are then presented with the CLI setup script.

Defaults or previously-entered values appear in brackets. To accept previously entered values, press **Enter**.

See the following guidelines:

- **Enter the IPv4 default gateway for the management interface**—If you set a manual IP address, enter either **data-interfaces** or the IP address of the gateway router. The **data-interfaces** setting sends outbound management traffic over the backplane to exit a data interface. This setting is useful if you do not have a separate Management network that can access the internet. Traffic originating on the Management interface includes license registration and database updates that require internet access. If you use **data-interfaces**, you can still use the FDM (or SSH) on the Management interface if you are directly-connected to the Management network, but for remote management for specific networks or hosts, you should add a static route using the **configure network static-routes** command. Note that the FDM management on data interfaces is not affected by this setting. If you use DHCP, the system uses the gateway provided by DHCP.
- **If your networking information has changed, you will need to reconnect**—If you are connected with SSH to the default IP address but you change the IP address at initial setup, you will be disconnected. Reconnect with the new IP address and password. Console connections are not affected. Note also that the DHCP server on Management will be disabled if you change the IP address.
- **Manage the device locally?**—Enter **yes** to use the FDM. A **no** answer means you intend to use the FMC to manage the device.

Example:

```
You must accept the EULA to continue.
Press <ENTER> to display the EULA:
End User License Agreement
[...]

Please enter 'YES' or press <ENTER> to AGREE to the EULA:

System initialization in progress. Please stand by.
You must configure the network to continue.
You must configure at least one of IPv4 or IPv6.
Do you want to configure IPv4? (y/n) [y]:
Do you want to configure IPv6? (y/n) [n]:
```

```

Configure IPv4 via DHCP or manually? (dhcp/manual) [manual]:
Enter an IPv4 address for the management interface [192.168.45.45]: 10.10.10.15
Enter an IPv4 netmask for the management interface [255.255.255.0]: 255.255.255.192
Enter the IPv4 default gateway for the management interface [data-interfaces]: 10.10.10.1
Enter a fully qualified hostname for this system [firepower]: ftd-1.cisco.com
Enter a comma-separated list of DNS servers or 'none' [208.67.222.222,208.67.220.220]:
Enter a comma-separated list of search domains or 'none' []:
If your networking information has changed, you will need to reconnect.
DHCP Server Disabled
The DHCP server has been disabled. You may re-enable with configure network ipv4
dhcp-server-enable

For HTTP Proxy configuration, run 'configure network http-proxy'

Manage the device locally? (yes/no) [yes]: yes

>

```

Step 4 Log into the FDM on the new Management IP address.

Complete the Initial Configuration Using the Setup Wizard

When you initially log into the FDM, you are taken through the device setup wizard to complete the initial system configuration.

Before you begin

Ensure that you connect a data interface to your gateway device, for example, a cable modem or router. For edge deployments, this would be your Internet-facing gateway. For data center deployments, this would be a back-bone router. Use the default “outside” interface for your model (see [Connect the Interfaces, on page 9](#) and [Default Configuration Prior to Initial Setup, on page 21](#)).

Then, connect your management computer to the “inside” interface for your hardware model. Alternatively, you can connect to the Management interface.

The Management interface does not need to be connected to a network. By default, the system obtains system licensing and database and other updates through the data interfaces, typically the outside interface, that connect to the internet. If you instead want to use a separate management network, you can connect the Management interface to a network and configure a separate management gateway after you complete initial setup.

To change the Management interface network settings if you cannot access the default IP address, see [\(Optional\) Change Management Network Settings at the CLI, on page 12](#).

Procedure

- Step 1** Log into the FDM.
- a) Assuming you did not go through initial configuration in the CLI, open the FDM at **https://ip-address**, where the address is one of the following.
 - If you are connected to the inside interface: **https://192.168.1.1**.
 - If you are connected to the Management interface: **https://192.168.45.45**.

b) Log in with the username **admin**. The default admin password is Admin123. .

Step 2 If this is the first time logging into the system, and you did not use the CLI setup wizard, you are prompted to read and accept the End User License Agreement and change the admin password.

You must complete these steps to continue.

Step 3 Configure the following options for the outside and management interfaces and click **Next**.

Caution Your settings are deployed to the device when you click **Next**. The interface will be named “outside” and it will be added to the “outside_zone” security zone. Ensure that your settings are correct. If you end up configuring an IP address on the outside interface that is on the same subnet as the inside interface, and you are connected to the FDM on the inside address, the wizard will hang when you click **Next**, because the address on the inside interface will be removed. To recover, see [What to Do if the Outside Subnet Conflicts with the Inside Subnet \(Setup Wizard Hangs at Step 1\)](#), on page 16.

Outside Interface

- **Configure IPv4**—The IPv4 address for the outside interface. You can use DHCP or manually enter a static IP address, subnet mask, and gateway. You can also select **Off** to not configure an IPv4 address. Do not configure an IP address on the same subnet as the default inside address (see [Default Configuration Prior to Initial Setup](#), on page 21), either statically or through DHCP.
- **Configure IPv6**—The IPv6 address for the outside interface. You can use DHCP or manually enter a static IP address, prefix, and gateway. You can also select **Off** to not configure an IPv6 address.

Management Interface

- **DNS Servers**—The DNS server for the system's management address. Enter one or more addresses of DNS servers for name resolution. The default is the OpenDNS public DNS servers. If you edit the fields and want to return to the default, click **Use OpenDNS** to reload the appropriate IP addresses into the fields. Your ISP might require that you use specific DNS servers. If after completing the wizard, you find that DNS resolution is not working, see [Troubleshooting DNS for the Management Interface](#), on page 305.
- **Firewall Hostname**—The hostname for the system's management address.

Step 4 Configure the system time settings and click **Next**.

- **Time Zone**—Select the time zone for the system.
- **NTP Time Server**—Select whether to use the default NTP servers or to manually enter the addresses of your NTP servers. You can add multiple servers to provide backups.

Step 5 Configure the smart licenses for the system.

You must have a smart license account to obtain and apply the licenses that the system requires. Initially, you can use the 90-day evaluation license and set up smart licensing later.

To register the device now, click the link to log into your Smart Software Manager account, generate a new token, and copy the token into the edit box.

If you do not want to register the device yet, select the evaluation mode option. The evaluation period last up to 90 days. To later register the device and obtain smart licenses, click **Device**, then click the link in the **Smart Licenses** group.

Step 6 Click **Finish**.**What to do next**

- If you want to use features covered by optional licenses, such as category-based URL filtering, intrusion inspection, or malware prevention, enable the required licenses. See [Enabling or Disabling Optional Licenses, on page 67](#).
- Connect the other data interfaces to distinct networks and configure the interfaces. For information on configuring interfaces, see [How to Add a Subnet, on page 57](#) and [Interfaces, on page 99](#).
- If you are managing the device through the inside interface, and you want to open CLI sessions through the inside interface, open the inside interface to SSH connections. See [Configuring the Management Access List, on page 281](#).
- Go through the use cases to learn how to use the product. See [Best Practices: Use Cases for FTD, on page 31](#).

What to Do if the Outside Subnet Conflicts with the Inside Subnet (Setup Wizard Hangs at Step 1)

If you connect to the FDM through the inside interface, you might find that the setup wizard hangs when you click **Next** during step 1, where you configure the outside interface. Note that normally it takes a while to complete this step, so hanging means that it continues for 10+ minutes. If you refresh the browser, you will see that you have lost the connection to the FDM. (If you connected through the management IP address, the wizard does not hang, but you might still have a problem as described in the symptoms below.)

The most likely reason this happens is that both the outside and inside interfaces were assigned addresses on the same subnet, which results in the inside interface losing its configuration.

The default configuration includes a static address on the inside interface, and a DHCP server, so that the device is functional and can pass traffic and support attached workstations immediately after you complete the setup wizard.

However, having a default inside address works only if you do not configure an address on the same subnet on the outside interface. This includes the situation where you attach to an ISP device that provides an address through DHCP to the outside address. Some ISPs use the same 192.168.1.0/24 subnet for their inside interface (which attaches to your outside interface) as FTD uses for the inside address.

To resolve this problem, you must change the IP address on the inside interface.

Symptoms for an inside/outside subnet conflict

Following are the symptoms that you have addresses on the same subnet on the inside and outside interfaces.




- During the device setup wizard, the wizard hangs when you click Next in step 1. Note that normally it takes a while to complete this step, so hanging means that it continues for 10+ minutes.
- If you are connected to the Console port, you would see the following message in the CLI. You will also get this message if you try to deploy the configuration (without subsequent change) from the FDM.

```
ERROR: Failed to apply IP address to interface GigabitEthernet1/1,  
as the network overlaps with interface GigabitEthernet1/2.
```


Two interfaces cannot be in the same subnet.

- If you get through setup, or exit it, the connection graphic will show no connection to any external services, such as the gateway, DNS and NTP servers, and Smart Licensing. The Deploy icon in the menu will also show that a deployment is needed.
- From the CLI, the **interface** and **dhcp** configurations are inconsistent for the inside and outside interfaces when viewed using the **show running-config** and **show startup-config** commands.

Procedure

- Step 1** If you were connected to the inside interface during device setup, complete the setup.
- Reconnect to the device by plugging into the Management port. If necessary, release and renew your workstation's DHCP address to get a new address on the management network (192.168.45.0/24). If necessary, configure a static address for your workstation in the 192.168.45.1-192.168.45.44 range.
 - Open the FDM at <https://192.168.45.45>.
 - You should see a prompt asking you to start your 90-day evaluation license. Select this option and click **Confirm**.
 - Choose **Device > System Settings > NTP**, configure the NTP servers, and click **Save**. If the default servers fit your requirements, you can skip this step.
 - Select **Profile** from the user icon drop-down list in the upper right of the menu, select the time zone for the device, and click **Save**.
- 
- If you do not want to use the evaluation license, choose **Device > Smart License > View Configuration**, click **Request Register**, then follow the instructions to register the device. See [Registering the Device, on page 66](#). (You can also enable any optional licenses you need at this time.)
- Step 2** Remove the DHCP server from the inside interface.
- Choose **Device > System Settings > DHCP Server**.
 - Click the **DHCP Servers** tab.
 - Mouse over the **Actions** column in the inside interface row and click the delete icon ()
- Step 3** Change the address on the inside interface.
- Select **Device**.
 - In the Interfaces group, click the link that indicates the number of enabled interfaces (for example, **3 Enabled**).
 - Mouse over the **Actions** column for the inside interface and click the edit icon ()
 - On the **IPv4 Address** tab, enter a static address on a unique subnet, for example, 192.168.2.1/24 or 192.168.46.1/24. Note that the default management address is 192.168.45.45/24, so do not use that subnet.
- You also have the option to use DHCP to obtain an address if you have a DHCP server already running on the inside network.
- Click **OK**.
- Step 4** (Optional.) Configure DHCP server on the inside address.

If you configure a static address for the inside interface, you can configure a DHCP server to provide addresses to workstations that attach to the inside network. This is a typical setup.

- a) Choose **Device > System Settings > DHCP Server**.
- b) Click the **DHCP Servers** tab.
- c) Click **+**.
- d) Select the option to enable the server and select the inside interface.
- e) For the address pool, enter a range on the same subnet as the inside address.

For example, if the inside address is 192.168.2.1/24, you might use 192.168.2.5-192.168.2.254. Do not include addresses that are statically assigned to nodes on the network. Consider leaving a few addresses outside the pool so you can assign static addresses when needed.

- f) Click **OK**.

Step 5 Click the **Deploy** button in the menu to deploy your changes.



Step 6 Click **Deploy Now**.

After deployment completes, the connection graphic should show green for the external services.

Configure the Wireless Access Point (ASA 5506W-X)

The ASA 5506W-X includes a Cisco Aironet 702i wireless access point integrated into the device. The wireless access point is disabled by default. Connect to the access point web interface so that you can enable the wireless radios and configure the SSID and security settings.

The access point connects internally over the GigabitEthernet1/9 interface. All Wi-Fi clients belong to the GigabitEthernet1/9 network. Your security policy determines how the Wi-Fi network can access any networks on other interfaces. The access point does not contain any external interfaces or switch ports.

The following procedure explains how to configure the access point. The procedure assumes that you completed the device setup wizard. If you instead manually configured the device, you might need to adjust the steps based on your configuration.

For more information, see the following manuals:

- For details about using the wireless LAN controller, see the [Cisco Wireless LAN Controller Software documentation](#).
- For details about the wireless access point hardware and software, see the [Cisco Aironet 700 Series documentation](#).

Before you begin

If you are unable to reach the access point, and the FTD device has the suggested configuration, and other networking issues are not found, then you may want to restore the access point default configuration. You must access the FTD CLI (connect to the console port, or configure SSH access). From the FTD CLI, enter the following commands.

```
> system support diagnostic-cli
Attaching to Diagnostic CLI ... Press 'Ctrl+a then d' to detach.
Type help or '?' for a list of available commands.

firepower> enable
Password: <press enter, by default, the password is blank>
firepower# hw-module module wlan recover configuration
```

If you need to troubleshoot the access point further, connect to the access point CLI using the **session wlan console** command.

Procedure

-
- Step 1** Configure and enable the wireless interface, GigabitEthernet1/9.
- Click **Device**, then click the link in the **Interfaces** group to open the list of interfaces.
 - Click the edit icon (🔗) for the GigabitEthernet1/9 interface.
 - Configure the following options.
 - **Interface Name**—Enter a name for the interface, for example, **wifi**.
 - **Status**—Click the slider to enable the interface.
 - **IPv4 Address**—Select **Static** for the address type, then enter an address and subnet mask. For example, 192.168.10.1/24.
 - Click **Save**.
- Step 2** Add the Wi-Fi interface to the same security zone as the inside interfaces.
- The device setup wizard puts the members of the **inside** bridge group in a security zone named **inside_zone**. The Wi-Fi interface needs to be in the same zone so that you can reach the access point web interface (made possible by the default Inside_Inside_Rule access rule).
- Click **Objects** in the menu, then select **Security Zones** from the table of contents.
 - Click the edit icon (🔗) for **inside_zone**.
 - Click + under **Interfaces** and select the **wifi** interface.
- Step 3** Verify that there is an access control rule to allow traffic between interfaces in the **inside_zone** security zone.
- The device setup wizard creates a rule to allow traffic to flow from the **inside_zone** to the **outside_zone**, which allows inside users to get to the Internet.
- The wizard also create a rule to allow traffic to flow between the **inside_zone** and **inside_zone**, so that internal hosts can reach each other.
- By adding the **wifi** interface to **inside_zone**, Wi-Fi users are also included in both of these rules, so that they can reach the Internet and other internal users.
- If you did not complete the wizard, these rules might not exist. Because the default action is to block all traffic, you must create these rules. The following procedure explains how to create a rule to enable traffic between the interfaces in the **inside_zone** security zone.
- Click **Policies** in the menu.
 - Click + above the **Access Control** table to add a rule.

- c) Configure at least the following options in the rule.
 - **Title**—Enter a name for the rule. For example, Inside_Inside.
 - **Action**—Either Allow or Trust.
 - **Source/Destination** > **Source Zones**—Select inside_zone.
 - **Source/Destination** > **Destination Zones**—Select inside_zone.
- d) Click **OK**.

Step 4 Configure the DHCP server on the wireless interface.

The DHCP server supplies IP addresses to devices that connect to the access point. It also supplies an address to the access point itself.

- a) Click **Device**.
- b) Click **System Settings** > **DHCP Server**.
- c) Click the **DHCP Servers** tab.
- d) Click + above the DHCP server table.
- e) Configure the following DHCP server properties.
 - **Enable DHCP Server**—Click the slider to enable the DHCP server.
 - **Interface**—Select the **wifi** interface.
 - **Address Pool**—Enter the address pool for DHCP clients. For example, if you used the example address for the wireless interface, the pool would be 192.168.10.2-192.168.10.254. The pool must be on the same subnet as the IP address for the interface, and it cannot include the address of the interface or the broadcast address.
- f) Click **OK**.

Step 5 Click the Deploy button in the menu, then click the **Deploy Now** button, to deploy your changes to the device.



Wait until the deployment finishes before you continue.

Step 6 Configure the wireless access point.

The wireless access point obtains its address from the DHCP pool defined for the wireless interface. It should get the first address in the pool. If you used the example addresses, this is 192.168.10.2. (Try the next address in the pool if the first one does not work.)

- a) Use a new browser window to go to the wireless access point IP address, for example, **http://192.168.10.2**.
The access point web interface should appear.
You must be on the inside network, or a network that can route to it, to open this address.
- b) Log in with the username **cisco** and password **Cisco**.
- c) On the left, click **Easy Setup** > **Network Configuration**.
- d) In the **Radio Configuration** area, for each of the **Radio 2.4GHz** and **Radio 5GHz** sections, set at least the following parameters and click **Apply** for each section.

- **SSID**—The Service Set Identifier. This is the name of the wireless network. Users will see this name when selecting a wireless network for their Wi-Fi connection.
- **Broadcast SSID in Beacon**—Select this option.
- **Universal Admin Mode: Disable**.
- **Security**—Select whichever security option you want to use.

- Step 7** While in the wireless access point web interface, enable the radios.
- On the left, click **Summary**, and then on the main page under **Network Interfaces**, click the link for the 2.4 GHz radio.
 - Click the **Settings** tab.
 - For the **Enable Radio** setting, click the **Enable** radio button, and then click **Apply** at the bottom of the page.
 - Repeat the process for the 5 GHz radio.

Default Configuration Prior to Initial Setup

Before you initially configure the FTD device using the local manager (FDM), the device includes the following default configuration.

For many models, this configuration assumes that you open the device manager through the inside interface, typically by plugging your computer directly into the interface, and use the DHCP server defined on the inside interface to supply your computer with an IP address. Alternatively, you can plug your computer into the Management interface and use DHCP to obtain an address. However, some models have different default configurations and management requirements. See the table below for details.



Note You can pre-configure many of these settings using the CLI setup ([\(\(Optional\) Change Management Network Settings at the CLI, on page 12\)](#)) before you perform setup using the wizard.

Default Configuration Settings

Setting	Default	Can be changed during initial configuration?
Password for admin user.	Admin123	Yes. You must change the default password.
Management IP address.	FTDv192.168.45.45	No.
Management gateway.	The data interfaces on the device. Typically the outside interface becomes the route to the Internet. This gateway works for from-the-device traffic only.	No.
DHCP server on the management interface.	Enabled with the address pool 192.168.45.46-192.168.45.254.	No.

Setting	Default	Can be changed during initial configuration?
DNS servers for the management interface.	The OpenDNS public DNS servers, 208.67.220.220 and 208.67.222.222.	Yes
Inside interface IP address.	192.168.1.1/24	No.
DHCP server for inside clients.	Running on the inside interface with the address pool 192.168.1.5 - 192.168.1.254.	No.
DHCP auto-configuration for inside clients. (Auto-configuration supplies clients with addresses for WINS and DNS servers.)	Enabled on outside interface.	Yes, but indirectly. If you configure a static IPv4 address for the outside interface, DHCP server auto-configuration is disabled.
Outside interface IP address.	Obtained through DHCP from Internet Service Provider (ISP) or upstream router.	Yes.

Default Interfaces by Device Model

You cannot select different inside and outside interfaces during initial configuration. To change the interface assignments after configuration, edit the interface and DHCP settings. You must remove an interface from the bridge group before you can configure it as a non-switched interface.

FTD device	Outside Interface	Inside Interface
ASA 5506-X ASA 5506H-X ASA 5506W-X	GigabitEthernet1/1	BV11, which contains all other data interfaces except the outside interface, and for the 5506W-X, the wireless interface GigabitEthernet1/9.
ASA 5508-X ASA 5516-X	GigabitEthernet1/1	GigabitEthernet1/2
ASA 5512-X ASA 5515-X ASA 5525-X ASA 5545-X ASA 5555-X	GigabitEthernet0/0	GigabitEthernet0/1

Configuration After Initial Setup

After you complete the setup wizard, the device configuration will include the following settings. The table shows whether a particular setting is something you explicitly chose or whether it was defined for you based on your other selections. Validate any "implied" configurations and edit them if they do not serve your needs.

Setting	Configuration	Explicit, implied, or default configuration
Password for admin user.	Whatever you entered.	Explicit.
Management IP address.	FTDv: 192.168.45.45	Default.
Management gateway.	The data interfaces on the device. Typically the outside interface becomes the route to the Internet. The management gateway works for from-the-device traffic only.	Default.
DHCP server on management interface.	Enabled with the address pool 192.168.45.46-192.168.45.254.	Default.
DNS servers for the management interface.	The OpenDNS public DNS servers, 208.67.220.220, 208.67.222.222, or whatever you entered. DNS servers obtained from DHCP are never used.	Explicit.
Management hostname.	firepower or whatever you entered.	Explicit.
Management access through data interfaces.	A data interface management access list rule allows HTTPS access through the inside interface. SSH connections are not allowed. Both IPv4 and IPv6 connections are allowed.	Implied.
System time.	The time zone and NTP servers you selected.	Explicit.
Smart license.	Either registered with a base license, or the evaluation period activated, whichever you selected. Subscription licenses are not enabled. Go to the smart licensing page to enable them.	Explicit.
Inside interface IP address.	192.168.1.1/24	Default.
DHCP server for inside clients.	Running on the inside interface with the address pool 192.168.1.5 - 192.168.1.254.	Default.
DHCP auto-configuration for inside clients. (Auto-configuration supplies clients with addresses for WINS and DNS servers.)	Enabled on outside interface if you use DHCP to obtain the outside interface IPv4 address. If you use static addressing, DHCP auto-configuration is disabled.	Explicit, but indirectly.
Data interface configuration.	<ul style="list-style-type: none"> ASA 5506-X—All data interfaces (such as GigabitEthernet1/2) except the outside interface are enabled and part of the inside bridge group. You can plug end points or switches into these ports and obtain addresses from the DHCP server for the inside interface. These interfaces are named inside_1, inside_2, and so forth. All other models—The outside and inside interfaces are the only ones configured and enabled. All other data interfaces are disabled. 	Default.

Setting	Configuration	Explicit, implied, or default configuration
Outside physical interface and IP address.	<p>The default outside port based on the device model. See Default Configuration Prior to Initial Setup, on page 21.</p> <p>The IP address is obtained by DHCP, or it is a static address as entered (IPv4, IPv6, or both).</p>	<p>Interface is Default.</p> <p>Addressing is Explicit.</p>
Static routes.	<p>If you configure a static IPv4 or IPv6 address for the outside interface, a static default route is configured for IPv4/IPv6 as appropriate, pointing to the gateway you defined for that address type. If you select DHCP, the default route is obtained from the DHCP server.</p> <p>Network objects are also created for the gateway and the "any" address, that is, 0.0.0.0/0 for IPv4, ::/0 for IPv6.</p>	Implied.
Security zones.	<p>inside_zone, containing the inside interface. For models that have an inside bridge group, the zone contains all members of the inside bridge group interface.</p> <p>outside_zone, containing the outside interface.</p> <p>(You can edit these zones to add other interfaces, or create your own zones.)</p>	Implied.
Access control policy.	<p>A rule trusting all traffic from the <code>inside_zone</code> to the <code>outside_zone</code>. This allows without inspection all traffic from users inside your network to get outside, and all return traffic for those connections.</p> <p>For models that have an inside bridge group, a second rule trusting all traffic between the interfaces in the <code>inside_zone</code>. This allows without inspection all traffic between users on your inside network.</p> <p>The default action for any other traffic is to block it. This prevents any traffic initiated from outside to enter your network.</p>	Implied.
NAT	<p>(Models that do not have an inside bridge group.) An interface dynamic PAT rule translates the source address for any IPv4 traffic destined to the outside interface to a unique port on the outside interface's IP address.</p> <p>(Models that have an inside bridge group.) For each member of the inside bridge group, an interface dynamic PAT rule translates the source address for any IPv4 traffic destined to the outside interface to a unique port on the outside interface's IP address. These appear in the NAT rule table and you can edit them later if desired.</p> <p>There are additional hidden PAT rules to enable HTTPS access through the inside interfaces, and routing through the data interfaces for the management address. These do not appear in the NAT table, but you will see them if you use the show nat command in the CLI.</p>	Implied.

Configuration Basics

The following topics explain the basic methods for configuring the device.

Configuring the Device

When you initially log into FDM, you are guided through a setup wizard to help you configure basic settings. Once you complete the wizard, use the following method to configure other features and to manage the device configuration.

If you have trouble distinguishing items visually, select a different color scheme in the user profile. Select **Profile** from the user icon drop-down menu in the upper right of the page.



Procedure

Step 1 Click **Device** to get to the **Device Summary**.

The dashboard shows a visual status for the device, including enabled interfaces and whether key settings are configured (colored green) or still need to be configured. For more information, see [Viewing Interface and Management Status, on page 28](#).

Above the status image is a summary of the device model, software version, VDB (System and Vulnerability Database) version, and the last time intrusion rules were updated.

Below the image are groups for the various features you can configure, with summaries of the configurations in each group, and actions you can take to manage the system configuration.

Step 2 Click the links in each group to configure the settings or perform the actions.

Following is a summary of the groups:

- **Interface**—You should have at least two data interfaces configured in addition to the management interface. See [Interfaces, on page 99](#).
- **Routing**—The routing configuration. You must define a default route. Other routes might be necessary depending on your configuration. See [Routing, on page 117](#).
- **Updates**—Geolocation, intrusion rule, and vulnerability database updates, and system software upgrades. Set up a regular update schedule to ensure that you have the latest database updates if you use those features. You can also go to this page if you need to download an update before the regularly schedule update occurs. See [Updating System Databases, on page 291](#).
- **System Settings**—This group includes a variety of settings. Some are basic settings that you would configure when you initially set up the device and then rarely change. See [System Settings, on page 281](#).
- **Smart License**—Shows the current state of the system licenses. You must install the appropriate licenses to use the system. Some features require additional licenses. See [Licensing the System, on page 63](#).
- **Backup and Restore**—Back up the system configuration or restore a previous backup. See [Backing Up and Restoring the System, on page 295](#).

- **Troubleshoot**—Generate a troubleshooting file at the request of the Cisco Technical Assistance Center. See [Creating a Troubleshooting File, on page 309](#).
- **Site-to-Site VPN**—The site-to-site virtual private network (VPN) connections between this device and remote devices. See [Managing Site-to-Site VPNs, on page 253](#).

Step 3 Click the **Deploy** button in the menu to deploy your changes.



Changes are not active on the device until you deploy them. See [Deploying Your Changes, on page 27](#).

What to do next

Click **Policies** in the main menu and configure the security policy for the system. You can also click **Objects** to configure the objects needed in those policies.

Configuring Security Policies

Use the security policies to implement your organization's acceptable use policy and to protect your network from intrusions and other threats.

Procedure

Step 1 Click **Policies**.

The Security Policies page shows the general flow of a connection through the system, and the order in which security policies are applied.

Step 2 Click the name of a policy and configure it.

You might not need to configure each policy type, although you must always have an access control policy. Following is a summary of the policies:

- **Identity**—If you want to correlate network activity to individual users, or control network access based on user or user group membership, use the identity policy to determine the user associated with a given source IP address. See [Configuring Identity Policies, on page 128](#).
- **NAT (Network Address Translation)**—Use the NAT policy to convert internal IP addresses to externally routeable addresses. See [Configure NAT, on page 173](#).
- **Access Control**—Use the access control policy to determine which connections are allowed on the network. You can filter by security zone, IP address, protocol, port, application, URL, user or user group. You also apply intrusion and file (malware) policies using access control rules. Use this policy to implement URL filtering. See [Configuring the Access Control Policy, on page 148](#).

Step 3 Click the **Deploy** button in the menu to deploy your changes.



Changes are not active on the device until you deploy them. See [Deploying Your Changes, on page 27](#).

Deploying Your Changes

When you update a policy or setting, the change is not immediately applied to the device. There is a two step process for making configuration changes:

1. Make your changes.
2. Deploy your changes.

This process gives you the opportunity to make a group of related changes without forcing you to run a device in a “partially configured” manner. In most cases, the deployment includes just your changes. However, if necessary, the system will reapply the entire configuration, which might be disruptive to your network. In addition, some changes require inspection engines to restart, with traffic dropping during the restart. Thus, consider deploying changes when potential disruptions will have the least impact.



Note If the deployment job fails, the system must roll back any partial changes to the previous configuration. Rollback includes clearing the data plane configuration and redeploying the previous version. This will disrupt traffic until the rollback completes.

After you complete the changes you want to make, use the following procedure to deploy them to the device.



Caution The FTD device drops traffic when the inspection engines are busy because of a software resource issue, or down because a configuration requires the engines to restart during configuration deployment. For detailed information on changes that require a restart, see [Configuration Changes that Restart Inspection Engines, on page 28](#).

Procedure

Step 1 Click the **Deploy Changes** icon in the upper right of the web page.

The icon is highlighted with a dot when there are undeployed changes.



The Deployment Summary page opens. The window shows a list of previous deployments with summary information on the changes (“modified objects”), when the deployment was initiated and completed, and the status of each deployment.

If the icon is not highlighted, you can still click it to see the results of previous deployment jobs.



Step 2 Click **Deploy Now**.

Configuration Changes that Restart Inspection Engines

Any of the following configurations or actions restart inspection engines when you deploy configuration changes.



Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations requires inspection engines to restart, which interrupts traffic inspection and drops traffic.

Deployment

Any deployment restarts the inspection engines.

System Updates

Installing a system update or patch that does not reboot the system and includes a binary change requires inspection engines to restart. Binary changes can include changes to inspection engines, a preprocessor, the vulnerability database (VDB), or a shared object rule. Note also that a patch that does not include a binary change can sometimes require a Snort restart.

Viewing Interface and Management Status

The Device Summary includes a graphical view of your device and select settings for the management address. To open the Device Summary, click **Device**.

Elements on this graphic change color based on the status of the element. Mousing over elements sometimes provides additional information. Use this graphic to monitor the following items.



Note The interface portion of the graphic, including interface status information, is also available on the **Interfaces** page and the **Monitoring > System** dashboard.

Interface Status

Mouse over a port to see its IP addresses, and enabled and link statuses. The IP addresses can be statically assigned or obtained using DHCP. Mousing over a Bridge Virtual Interface (BVI) also shows the list of member interfaces.

Interface ports use the following color coding:

- Green—The interface is configured, enabled, and the link is up.
- Gray—The interface is not enabled.
- Orange/Red—The interface is configured and enabled, but the link is down. If the interface is wired, this is an error condition that needs correction. If the interface is not wired, this is the expected status.

Inside, Outside Network Connections

The graphic indicates which port is connected to the outside (or upstream) and inside networks, under the following conditions.

- Inside Network—The port for the inside network is shown for the interface named “inside” only. If there are additional inside networks, they are not shown. If you do not name any interface “inside,” no port is marked as the inside port.
- Outside Network—The port for the outside network is shown for the interface named “outside” only. As with the inside network, this name is required, or no port is marked as the outside port.

Management Setting Status

The graphic shows whether the gateway, DNS servers, NTP servers, and Smart Licensing are configured for the management address, and whether those settings are functioning correctly.

Green indicates that the feature is configured and functioning correctly, gray indicates that it is not configured or not functioning correctly. For example, the DNS box is gray if the servers cannot be reached. Mouse over the elements to see more information.

If you find problems, correct them as follows:

- Management port and gateway—Select **System Settings** > **Management Interface**.
- DNS servers—Select **System Settings** > **DNS Server**.
- NTP servers—Select **System Settings** > **NTP**. Also see [Troubleshooting NTP, on page 303](#).
- Smart License—Click the **View Configuration** link in the Smart License group.

Viewing System Task Status

System tasks include actions that occur without your direct involvement, such as retrieving and applying various database updates. You can view a list of these tasks and their status to verify that these system tasks are completing successfully.

Procedure

Step 1 Click the **Task List** button in the main menu.




The task list opens, displaying the status and details of system tasks.

Step 2 Evaluate the task status.

If you find a persistent problem, you might need to fix the device configuration. For example, a persistent failure to obtain database updates could indicate that there is no path to the Internet for the device's management IP address. You might need to contact the Cisco Technical Assistance Center (TAC) for some issues as indicted in the task descriptions.

You can do the following with the task list:

- Click the **Success** or **Failures** buttons to filter the list based on these statuses.

- Click the delete icon () for a task to remove it from the list.
 - Click **Remove All Completed Tasks** to empty the list of all tasks that are not in progress.
-



CHAPTER 2

Best Practices: Use Cases for FTD

The following topics explain some common tasks you might want to accomplish with FTD using the FDM. These use cases assume that you completed the device configuration wizard and that you retained this initial configuration. Even if you modified the initial configuration, you should be able to use these examples to understand how to use the product.

- [How to Configure the Device in FDM, on page 31](#)
- [How to Gain Insight Into Your Network Traffic, on page 36](#)
- [How to Block Threats, on page 42](#)
- [How to Block Malware, on page 46](#)
- [How to Implement an Acceptable Use Policy \(URL Filtering\), on page 49](#)
- [How to Control Application Usage, on page 53](#)
- [How to Add a Subnet, on page 57](#)
- [More Examples, on page 62](#)

How to Configure the Device in FDM

After you complete the setup wizard, you should have a functioning device with a few basic policies in place:

- (Except for ASA 5506-X and.) An outside and an inside interface. No other data interfaces are configured.
- (ASA 5506-X, only.) An outside interface, and an inside bridge group that includes all other data interfaces.
- Security zones for the inside and outside interfaces.
- An access rule trusting all inside to outside traffic.
- An interface NAT rule that translates all inside to outside traffic to unique ports on the IP address of the outside interface.
- A DHCP server running on the inside interface or bridge group.

The following steps provide an overview of additional features you might want to configure. Please click the help button (?) on a page to get detailed information about each step.

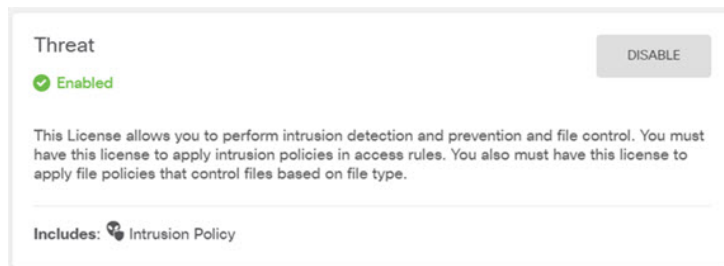
Procedure

Step 1 Choose **Device**, then click **View Configuration** in the **Smart License** group.

Click **Enable** for each of the optional licenses you want to use: Threat, Malware, URL. Read the explanation of each license if you are unsure of whether you need it.

If you have not registered, you can do so from this page. Click **Request Register** and follow the instructions. Please register before the evaluation license expires.

For example, an enabled Threat license should look like the following:



Step 2 If you wired other interfaces, choose **Device**, then click the link in the **Interfaces** summary.

- Because the ASA 5506-X and comes pre-configured with a bridge group containing all non-outside data interfaces, there is no need to configure these interfaces. If you want to break apart the bridge group, you can edit it to remove the interfaces you want to treat separately. Then you can configure those interfaces as hosting separate networks.

For other models, you can create a bridge group for the other interfaces, or configure separate networks, or some combination of both.

Click the edit icon (🔗) for each interface to define the IP address and other settings.

The following example configures an interface to be used as a “demilitarized zone” (DMZ), where you place publically-accessible assets such as your web server. Click **Save** when you are finished.

Edit Physical Interface

Interface Name Status

dmz

Description

[IPv4 Address](#) [IPv6 Address](#) [Advanced Options](#)

Type

Static

IP Address and Subnet Mask

192.168.6.1 / 24

e.g. 192.168.5.15/17 or 192.168.5.15/255.255.128.0

- Step 3** If you configured new interfaces, choose **Objects**, then select **Security Zones** from the table of contents.
- Edit or create new zones as appropriate. Each interface must belong to a zone, because you configure policies based on security zones, not interfaces. You cannot put the interfaces in zones when configuring them, so you must always edit the zone objects after creating new interfaces or changing the purpose of existing interfaces.
- The following example shows how to create a new dmz-zone for the dmz interface.

Add Security Zone

Name

dmz-zone

Description

Interfaces

dmz

- Step 4** If you want internal clients to use DHCP to obtain an IP address from the device, choose **Device**, then **System Settings > DHCP Server**. Select the **DHCP Servers** tab.

There is already a DHCP server configured for the inside interface, but you can edit the address pool or even delete it. If you configured other inside interfaces, it is very typical to set up a DHCP server on those interfaces. Click + to configure the server and address pool for each inside interface.

You can also fine-tune the WINS and DNS list supplied to clients on the **Configuration** tab.

The following example shows how to set up a DHCP server on the inside2 interface with the address pool 192.168.4.50-192.168.4.240.

Add Server

Enabled DHCP Server

Interface
inside2

Address Pool
192.168.4.50-192.168.4.240
e.g. 192.168.45.46-192.168.45.254

Step 5 Choose **Device**, then click **View Configuration** (or **Create First Static Route**) in the **Routing** group and configure a default route.

The default route normally points to the upstream or ISP router that resides off the outside interface. A default IPv4 route is for any-ipv4 (0.0.0.0/0), whereas a default IPv6 route is for any-ipv6 (:::0/0). Create routes for each IP version you use. If you use DHCP to obtain an address for the outside interface, you might already have the default routes that you need.

The routes you define on this page are for the data interfaces only. They do not impact the management interface. Set the management gateway on **System Settings > Management Interface**.

The following example shows a default route for IPv4. In this example, isp-gateway is a network object that identifies the IP address of the ISP gateway (you must obtain the address from your ISP). You can create this object by clicking **Create New Network** at the bottom of the **Gateway** drop-down list.

Add Static Route

Protocol
 IPv4 IPv6

Gateway
isp-gateway

Interface
outside

Metric
1

Networks
+
any-ipv4

Step 6 Choose **Policies** and configure the security policies for the network.


The device setup wizard enables traffic flow between the inside-zone and outside-zone, and interface NAT for all interfaces when going to the outside interface. Even if you configure new interfaces, if you add them to the inside-zone object, the access control rule automatically applies to them.

However, if you have multiple inside interfaces, you need an access control rule to allow traffic flow from inside-zone to inside-zone. If you add other security zones, you need rules to allow traffic to and from those zones. These would be your minimum changes.

In addition, you can configure other policies to provide additional services, and fine-tune NAT and access rules to get the results that your organization requires. You can configure the following policies:

- **Identity**—If you want to correlate network activity to individual users, or control network access based on user or user group membership, use the identity policy to determine the user associated with a given source IP address.
- **NAT (Network Address Translation)**—Use the NAT policy to convert internal IP addresses to externally routeable addresses.
- **Access Control**—Use the access control policy to determine which connections are allowed on the network. You can filter by security zone, IP address, protocol, port, application, URL, user or user group. You also apply intrusion and file (malware) policies using access control rules. Use this policy to implement URL filtering.

The following example shows how to allow traffic between the inside-zone and dmz-zone in the access control policy. In this example, no options are set on any of the other tabs except for **Logging**, where **At End of Connection** is selected.



Step 7 Commit your changes.

- a) Click the **Deploy Changes** icon in the upper right of the web page.



- b) Click the **Deploy Now** button.

Wait for deployment to finish. The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.

How to Gain Insight Into Your Network Traffic

After completing initial device setup, you have an access control policy that allows all inside traffic access to the Internet or other upstream network, and a default action to block all other traffic. Before you create additional access control rules, you might find it beneficial to gain insight into the traffic that is actually occurring on your network.

You can use the monitoring capabilities of the FDM to analyze network traffic. FDM reporting helps you answer the following questions:

- What is my network being used for?
- Who is using the network the most?
- Where are my users going?
- What devices are they using?
- What access control rules (policies) are being hit the most?

The initial access rule can provide some insight into traffic, including policies, destinations, and security zones. But to obtain user information, you need to configure an identity policy that requires users to authenticate (identify) themselves. To obtain information on applications used on the network, you need to make some additional adjustments.

The following procedure explains how to set up the FTD device to monitor traffic and provides an overview of the end-to-end process of configuring and monitoring policies.



Note This procedure does not provide insight into the web site categories and reputations of sites visited by users, so you cannot see meaningful information in the web categories dashboard. You must implement category-based URL filtering, and enable the URL license, to obtain category and reputation data. If you just want to obtain this information, you can add a new access control rule that allows access to an acceptable category, such as Financial Services, and make it the first rule in the access control policy. For details on implementing URL filtering, see [How to Implement an Acceptable Use Policy \(URL Filtering\)](#), on page 49.

Procedure

Step 1

To gain insight into user behavior, you need to configure an identity policy to ensure that the user associated with a connection is identified.

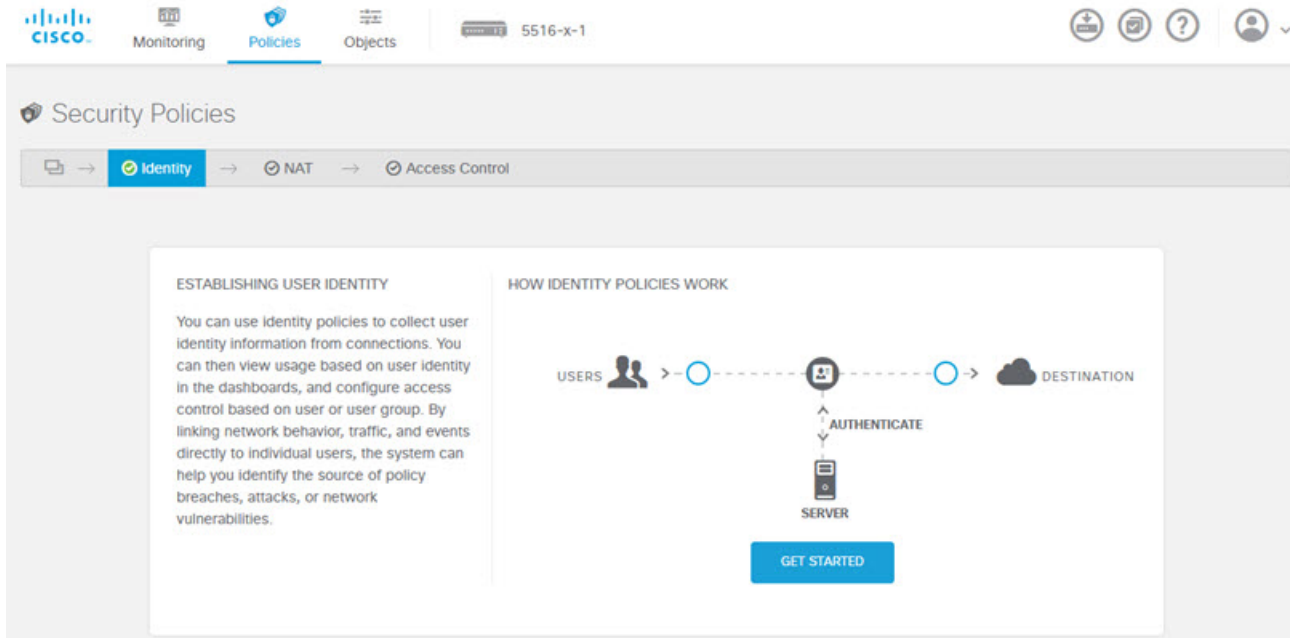
By enabling the identity policy, you can collect information about who is using the network, and what resources they are using. This information is available in the User monitoring dashboard. User information is also available for connection events shown in Event Viewer.

Users are authenticated only when they use a web browser for HTTP connections.

If a user fails to authenticate, the user is not prevented from making web connections. This just means that you do not have user identity information for the connections. If you want, you can create an access control rule to drop traffic for Failed Authentication users.

- a) Click **Policies** in the main menu, then click **Identity**.

The identity policy is initially disabled. The identity policy uses your Active Directory server to authenticate users and associate them with the IP address of the workstation they are using. Subsequently, the system will identify traffic for that IP address as being the user's traffic.



- b) Click the **Get Started** button to start the wizard to configure the needed elements.
 c) Identify your Active Directory server.

Fill in the following information.

- **Name**—A name for the directory realm.
- **Type**—The type of directory server. Active Directory is the only supported type, and you cannot change this field.
- **Directory Username, Directory Password**—The distinguished username and password for a user with appropriate rights to the user information you want to retrieve. For Active Directory, the user does not need elevated privileges. You can specify any user in the domain. The username must be fully qualified; for example, Administrator@example.com (not simply Administrator).
- **Base DN**—The directory tree for searching or querying user and group information, that is, the common parent for users and groups. For example, dc=example,dc=com. For information on finding the base DN, see [Determining the Directory Base DN, on page 126](#).
- **AD Primary Domain**—The fully qualified Active Directory domain name that the device should join. For example, example.com.
- **Hostname/IP Address**—The hostname or IP address of the directory server. If you use an encrypted connection to the server, you must enter the fully-qualified domain name, not the IP address.
- **Port**—The port number used for communications with the server. The default is 389. Use port 636 if you select LDAPS as the encryption method.

- **Encryption**—To use an encrypted connection for downloading user and group information, select the desired method, **STARTTLS** or **LDAPS**. The default is **None**, which means that user and group information is downloaded in clear text.
 - **STARTTLS** negotiates the encryption method, and uses the strongest method supported by the directory server. Use port 389.
 - **LDAPS** requires LDAP over SSL. Use port 636.
- **SSL Certificate**—If you select an encryption method, upload a CA certificate to enable a trusted connection between the system and the directory server. If you are using a certificate to authenticate, the name of the server in the certificate must match the server Hostname / IP Address. For example, if you use 10.10.10.250 as the IP address but ad.example.com in the certificate, the connection fails.

Example:

For example, the following image shows how to create an unencrypted connection for the ad.example.com server. The primary domain is example.com, and the directory username is Administrator@ad.example.com. All user and group information is under the Distinguished Name (DN) ou=user,dc=example,dc=com.

The screenshot shows a configuration window for a Directory Server. The fields are as follows:

Field	Value
Name	AD
Type	Active Directory (AD)
Directory Username	Administrator@ad.example.com
Directory Password
Base DN	ou=user,dc=example,dc=com
AD Primary Domain	example.com
Hostname / IP Address	ad.example.com
Port	389
Encryption	NONE
SSL Certificate	UPLOAD button, No certificates uploaded yet.

- Click **Next**.
- Configure the active authentication captive portal.

The simplest option is to leave all fields as is and click **Save**. You would configure the default port for active authentication, and users get a self-signed certificate that they need to trust in order to provide their username and password. Tell users to expect this and that they should accept the certificate.

However, you would ideally upload a certificate that their browsers already trust. If you have one, fill in the following fields to use it.

- **Server Certificate**—The CA certificate to present to users during active authentication. The certificate must be an X509 certificate in PEM or DER format. Paste in the certificate, or click **Upload Certificate** and select the certificate file. The default is to present a self-signed certificate during user authentication.
- **Certificate Key**—The key for the server certificate. Paste in the key, or click **Upload Key** and select the key file.
- **Port**—The captive portal port. The default is 885 (TCP). If you configure a different port, it must be in the range 1025-65535.

f) Click **Save**.

This completes the setup wizard. Now, create an identity rule to require active authentication.

g) Click the **Create Identity Rule** button, or the + button.

h) Fill in the identity rule properties.

Assuming you want to require everyone to authenticate, you could use the following settings:

- **Name**—Anything you choose, for example, `Require_Authentication`.
- **User Authentication**—**Active** should already be selected; keep it.
- **Type**—Select **HTTP Negotiate**. This allows the browser and directory server to negotiate the strongest authentication protocol, in order, NTLM, then HTTP basic.

Note For the HTTP Basic, HTTP Response Page, and NTLM authentication methods, the user is redirected to the captive portal using the IP address of the interface. However, for HTTP Negotiate, the user is redirected using the fully-qualified DNS name *firewall-hostname.AD-domain-name*. If you want to use HTTP Negotiate, you must also update your DNS server to map this name to the IP addresses of all inside interfaces where you are requiring active authentication. Otherwise, the redirection cannot complete, and users cannot authenticate. If you cannot, or do not want to, update the DNS server, select one of the other authentication methods.

- **Source/Destination**—Leave all fields to default to Any.

You can constrain the policy as you see fit to a more limited set of traffic. However, active authentication will only be attempted for HTTP traffic, so it does not matter that non-HTTP traffic matches the source/destination criteria. For more details about identity policy properties, see [Configure Identity Rules, on page 131](#).

- i) Click **OK** to add the rule.

If you look in the upper right of the window, you can see that the **Deploy** icon button now has a dot, which indicates that there are undeployed changes. Making changes in the user interface is not sufficient for getting the changes configured on the device, you must deploy changes. Thus, you can make a set of related changes before you deploy them, so that you do not face the potential problems of having a partially-configured set of changes running on the device. You will deploy changes later in this procedure.



Step 2 Change the action on the Inside_Outside_Rule access control rule to **Allow**.

The Inside_Outside_Rule access rule is created as a trust rule. However, trusted traffic is not inspected, so the system cannot learn about some of the characteristics of trusted traffic, such as application, when the traffic matching criteria does not include application or other conditions besides zone, IP address, and port. If you change the rule to allow rather than trust traffic, the system fully inspects the traffic.

Note (ASA 5506-X .) Also consider changing the Inside_Inside_Rule from Trust to Allow. This rule covers traffic going between the inside interfaces.


- a) Click **Access Control** on the **Policies** page.
- b) Hover over the **Actions** cell on the right side of the Inside_Outside_Rule row to expose the edit and delete icons, and click the edit icon (🔧) to open the rule.
- c) Select **Allow** for the **Action**.

- d) Click **OK** to save the change.

Step 3 Enable logging on the access control policy default action.

Dashboards contain information about connections only if the connection matches an access control rule that enables connection logging. The Inside_Outside_Rule enables logging, but the default action has logging disabled. Thus, dashboards show information for the Inside_Outside_Rule only, and do not reflect connections that do not match any rules.

- a) Click anywhere in the default action at the bottom of the access control policy page.



- b) Select **Select Log Action** > **At Beginning and End of Connection**.
- c) Click **OK**.

Step 4 Set an update schedule for the vulnerability database (VDB).

Cisco regularly releases updates to the VDB, which includes the application detectors that can identify the application used in a connection. You should update the VDB on a regular basis. You can either manually download updates, or you can set up a regular schedule. The following procedure shows how to set up a schedule. By default, VDB updates are disabled, so you need to take action to get VDB updates.

- a) Click **Device**.
- b) Click **View Configuration** in the Updates group.

Updates

[View Configuration](#) >

- c) Click **Configure** in the VDB group.

VDB 265.0

Configure
Set recurring VDB updates

UPDATE NOW i

- d) Define the update schedule.

Choose a time and frequency that will not be disruptive to your network. Also, please understand that the system will do an automatic deployment after downloading the update. This is necessary to activate the new detectors. Thus, any configuration changes that you have made and saved but have not yet deployed will also be deployed.

For example, the following schedule updates the VDB once a week on Sunday at 12:00 AM (using the 24-hour clock notation).

Set recurring VDB Update

Frequency

Weekly

Days of Week

Sundays x at 00 : 00

(-07:00) America/Los_Angeles

- e) Click **Save**.

- Step 5** Commit your changes.
- Click the **Deploy Changes** icon in the upper right of the web page.



- Click the **Deploy Now** button and wait for deployment to finish.

The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.

Deployment Summary ?

DEPLOY NOW
You have successfully deployed.

Deployment History

Modified Objects	Initiated	Completed	Status
> AccessPolicy	11 May 2016	11 May 2016	✔ Deployed
> AccessRule	01:24:35 PM	01:27:06 PM	
> ActiveDirectoryRealm			
> IdentityPolicy			
> IdentityRule			

What to do next

At this point, the monitoring dashboards and events should start showing information about users and applications. You can evaluate this information for undesirable patterns and develop new access rules to constrain unacceptable use.

If you want to start collecting information about intrusions and malware, you need to enable intrusion and file policies on one or more access rule. You also need to enable the licenses for those features.

If you want to start collecting information about URL categories, you must implement URL filtering.

How to Block Threats

You can implement next generation Intrusion Prevention System (IPS) filtering by adding intrusion policies to your access control rules. Intrusion policies analyze network traffic, comparing the traffic contents against known threats. If a connection matches a threat you are monitoring, the system drops the connection, thus preventing the attack.

All other traffic handling occurs before network traffic is examined for intrusions. By associating an intrusion policy with an access control rule, you are telling the system that before it passes traffic that matches the access control rule's conditions, you first want to inspect the traffic with an intrusion policy.

You can configure intrusion policies on rules that **allow** traffic only. Inspection is not performed on rules set to **trust** or **block** traffic. In addition, you can configure an intrusion policy as part of the default action if the default action is **allow**.

The intrusion policies are designed by the Cisco Talos Intelligence Group (Talos), who set the intrusion and preprocessor rule states and advanced settings.

Procedure

Step 1 If you have not already done so, enable the Threat license.

You must enable the Threat license to use intrusion policies. If you are currently using the evaluation license, you are enabling an evaluation version of the license. If you have registered the device, you must purchase the required license and add it to your Smart Software Manager account on Cisco.com.

- a) Click **Device**.
- b) Click **View Configuration** in the Smart License group.



- c) Click **Enable** in the **Threat** group.


The system registers the license with your account, or activates the evaluation license, as appropriate. The group should indicate that the license is enabled, and the button changes to a Disable button.



Step 2 Select an intrusion policy for one or more access rules.

Determine which rules cover traffic that should be scanned for threats. For this example, we will add intrusion inspection to the Inside_Outside_Rule. For ASA 5506-X models, you might also want to add it to the Inside_Inside_Rule.

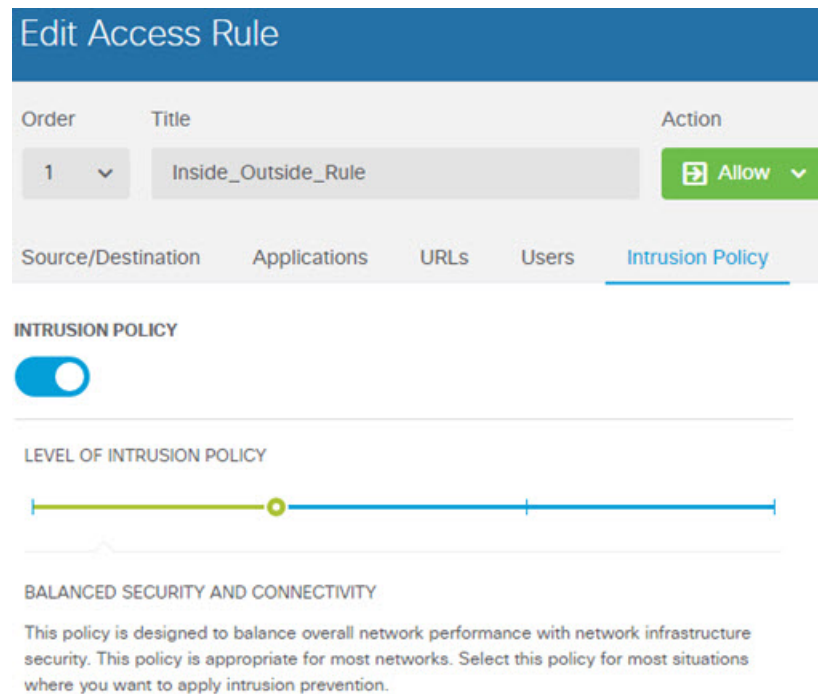
- a) Click **Policies** in the main menu.
Ensure that the **Access Control** policy is displayed.
- b) Hover over the **Actions** cell on the right side of the Inside_Outside_Rule row to expose the edit and delete icons, and click the edit icon (🔧) to open the rule.
- c) If you have not already done so, select **Allow** for the **Action**.



- d) Click the **Intrusion Policy** tab.
- e) Click the **Intrusion Policy** toggle to enable it, then select the intrusion policy.

The **Balanced Security and Connectivity** policy is appropriate for most networks. It provides a good intrusion defense without being overly aggressive, which has the potential of dropping traffic that you might not want to be dropped. If you determine that too much traffic is getting dropped, you can ease up on intrusion inspection by selecting the **Connectivity over Security** policy.

If you need to be aggressive about security, try the **Security over Connectivity** policy. The **Maximum Detection** policy offers even more emphasis on network infrastructure security with the potential for even greater operational impact.




f) Click **OK** to save the change.

Step 3 Set an update schedule for the intrusion rule database.

Cisco regularly releases updates to the intrusion rule database, which is used by intrusion policies to determine whether connections should be dropped. You should update the rule database on a regular basis. You can either manually download updates, or you can set up a regular schedule. The following procedure shows how to set up a schedule. By default, database updates are disabled, so you need to take action to get updated rules.

- a) Click **Device**.
- b) Click **View Configuration** in the Updates group.

Updates



c) Click **Configure** in the Rule group.

Rule 2016-03-28-001-vrt

Configure

Set recurring Rule updates

UPDATE NOW



- d) Define the update schedule.

Choose a time and frequency that will not be disruptive to your network. Also, please understand that the system will do an automatic deployment after downloading the update. This is necessary to activate the new rules. Thus, any configuration changes that you have made and saved but have not yet deployed will also be deployed.

For example, the following schedule updates the rule database once a week on Monday at 12:00 AM (using the 24-hour clock notation).

Set recurring Rule Update

Frequency

Weekly

Days of Week

Mondays ×

Time

at 00 : 00

(-07:00) America/Los_Angeles

- e) Click **Save**.

Step 4

Commit your changes.

- a) Click the **Deploy Changes** icon in the upper right of the web page.



- b) Click the **Deploy Now** button.

Wait for deployment to finish. The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.

What to do next

At this point, the monitoring dashboards and events should start showing information about attackers, targets, and threats, if any intrusions are identified. You can evaluate this information to determine if your network needs more security precautions, or if you need to reduce the level of intrusion policy you are using.

How to Block Malware

Users are continually at risk of obtaining malicious software, or *malware*, from Internet sites or other communication methods, such as e-mail. Even trusted web sites can be hijacked to serve malware to unsuspecting users. Web pages can contain objects coming from different sources. These objects can include images, executables, Javascript, advertisements, and so forth. Compromised web sites often incorporate objects hosted on external sources. Real security means looking at each object individually, not just the initial request.

Use file policies to detect malware using malware defense. You can also use file policies to perform file control, which allows control over all files of a specific type regardless of whether the files contain malware.

Malware defense uses the AMP Cloud to retrieve dispositions for possible malware detected in network traffic. The management interface must have a path to the Internet to reach the AMP Cloud and perform malware lookups. When the device detects an eligible file, it uses the file's SHA-256 hash value to query the AMP Cloud for the file's disposition. The possible disposition can be **clean**, **malware**, or **unknown** (no clear verdict). If the AMP Cloud is unreachable, the disposition is **unknown**.

By associating a file policy with an access control rule, you are telling the system that before it passes traffic that matches the access control rule's conditions, you first want to inspect any files in the connection.

You can configure file policies on rules that **allow** traffic only. Inspection is not performed on rules set to **trust** or **block** traffic.

Procedure

Step 1

If you have not already done so, enable the Malware and Threat licenses.

You must enable the Malware to use file policies in addition to the Threat license, which is required for intrusion policies. If you are currently using the evaluation license, you are enabling an evaluation version of the licenses. If you have registered the device, you must purchase the required licenses and add them to your Smart Software Manager account on Cisco.com.

- a) Click **Device**.
- b) Click **View Configuration** in the Smart License group.

Smart License

Registered

[View Configuration](#) >

- c) Click **Enable** in the **Malware** group, and if not already enabled, the **Threat** group.

The system registers the license with your account, or activates the evaluation license, as appropriate. The group should indicate that the license is enabled, and the button changes to a Disable button.

Malware

✓ Enabled

DISABLE

Step 2 Select a file policy for one or more access rules.

Determine which rules cover traffic that should be scanned for malware. For this example, we will add file inspection to the Inside_Outside_Rule. For ASA 5506-X models, you might also want to add it to the Inside_Inside_Rule.

- Click **Policies** in the main menu.
Ensure that the **Access Control** policy is displayed.
- Hover over the **Actions** cell on the right side of the Inside_Outside_Rule row to expose the edit and delete icons, and click the edit icon (🔗) to open the rule.
- If you have not already done so, select **Allow** for the **Action**.

Order	Title	Action
1	Inside_Outside_Rule	Allow

- Click the **File Policy** tab.
- Click the file policy you want to use.

Your main choice is between **Block Malware All**, which drops any files that are considered malware, or **Cloud Lookup All**, which queries the AMP Cloud to determine the file's disposition, but does no blocking. If you want to first see how files are being evaluated, use cloud lookup. You can switch to the blocking policy later if you are satisfied with how files are being evaluated.

There are other policies available that block malware. These policies are coupled with file control, blocking the upload of Microsoft Office, or Office and PDF, documents. That is, these policies prevent users from sending these file types to other networks in addition to blocking malware. You can select these policies if they fit your needs.

For this example, select **Block Malware All**.

The screenshot shows the configuration page for rule '1' titled 'Editing Rule Inside_Outside_Rule'. The 'Name' field contains 'Inside_Outside_Rule'. The 'Logging' is set to 'ON', 'Time Range' is 'None', and 'Rule Enabled' is checked. The 'File Policy' dropdown is set to 'Block Malware All'. Below the main configuration area, there is a navigation bar with tabs for 'All', 'Zones', 'Networks', 'Ports', 'Applications', 'Users', 'URLs', 'Dynamic Attributes', and 'VLAN Tags', with 'All' currently selected.

- f) Click the **Logging** tab and verify that **Log Files** under File Events is selected.

By default, file logging is enabled whenever you select a file policy. You must enable file logging to get file and malware information in events and dashboards.

FILE EVENTS

Log Files

- g) Click **OK** to save the change.

Step 3

Commit your changes.

- a) Click the **Deploy Changes** icon in the upper right of the web page.



- b) Click the **Deploy Now** button.

Wait for deployment to finish. The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.

What to do next

At this point, the monitoring dashboards and events should start showing information about file types and file and malware events, if any files or malware are transmitted. You can evaluate this information to determine if your network needs more security precautions related to file transmissions.

How to Implement an Acceptable Use Policy (URL Filtering)

You might have an acceptable use policy for your network. Acceptable use policies differentiate between network activity that is appropriate in your organization and activity that is considered inappropriate. These policies are typically focused on Internet usage, and are geared towards maintaining productivity, avoiding legal liabilities (for example, maintaining a non-hostile workplace), and in general controlling web traffic.

You can use URL filtering to define an acceptable use policy with access policies. You can filter on broad categories, such as Gambling, so that you do not need to identify every individual web site that should be blocked. For category matches, you can also specify the relative reputation of sites to allow or block. If a user attempts to browse to any URL with that category and reputation combination, the session is blocked.

Using category and reputation data also simplifies policy creation and administration. It grants you assurance that the system will control web traffic as expected. Finally, because Cisco's threat intelligence is continually updated with new URLs, as well as new categories and risks for existing URLs, you can ensure that the system uses up-to-date information to filter requested URLs. Malicious sites that represent security threats such as malware, spam, botnets, and phishing may appear and disappear faster than you can update and deploy new policies.

The following procedure explains how to implement an acceptable use policy using URL filtering. For purposes of this example, we will block sites of any reputation in several categories, risky Social Networking sites, and an unclassified site, badsite.example.com.

Procedure

Step 1 If you have not already done so, enable the **URL** license.

You must enable the URL license to use URL category and reputation information, or to see the information in dashboards and events. If you are currently using the evaluation license, you are enabling an evaluation version of the license. If you have registered the device, you must purchase the required license and add it to your Smart Software Manager account on Cisco.com.

- a) Click **Device**.
- b) Click **View Configuration** in the Smart License group.



- c) Click **Enable** in the **URL License** group.

The system registers the license with your account, or activates the evaluation license, as appropriate. The group should indicate that the license is enabled, and the button changes to a Disable button.



Step 2 Create a URL filtering access control rule.

You might want to first see the categories for sites your users are visiting before making a blocking rule. If that is the case, you can create a rule with the Allow action for an acceptable category, such as Financial Services. Because all web connections must be inspected to determine if the URL belongs to this category, you would get category information even for non-Financial Services sites.

But there are probably URL categories that you already know you want to block. A blocking policy also forces inspection, so you get category information on connections to unblocked categories, not just the blocked categories.

- a) Click **Policies** in the main menu.

Ensure that the **Access Control** policy is displayed.

- b) Click + to add a new rule.

- c) Configure the order, title, and action.

- **Order**—The default is to add new rules to the end of the access control policy. However, you must place this rule ahead of (above) any rule that would match the same Source/Destination and other criteria, or the rule will never be matched (a connection matches one rule only, and that is the first rule it matches in the table). For this rule, we will use the same Source/Destination as the `Inside_Outside_Rule` created during initial device configuration. You might have created other rules as well. To maximize access control efficiency, it is best to have specific rules early, to ensure the quickest decision on whether a connection is allowed or dropped. For the purposes of this example, select **1** as the rule order.


- **Title**—Give the rule a meaningful name, such as `Block_Web_Sites`.

- **Action**—Select **Block**.


Order	Title	Action
1	Block_Web_Sites	Block

- d) On the **Source/Destination** tab, click + for **Source > Zones**, select `inside_zone`, then click **OK** in the zones dialog box.

Adding any of the criteria works the same way. Clicking + opens a little dialog box, where you click the items you want to add. You can click multiple items, and clicking a selected item de-selects it; the check marks indicate the selected items. But nothing is added to the policy until you click the **OK** button; simply selecting the items is not sufficient.




- e) Using the same technique, select **outside_zone** for **Destination > Zones**.




- f) Click the **URLs** tab.
- g) Click the + for **Categories**, and select the categories you want to fully or partially block.

For purposes of this example, select Adult and Pornography, Bot Nets, Confirmed SPAM Sources, and Social Network. There are additional categories that you would most likely want to block.



- h) To implement reputation-sensitive blocking for the Social Network category, click **Reputation: Risk Any** for that category, deselect **Any**, then move the slider to **Benign sites with security risks**. Click away from the slider to close it.



The left of the reputation slider indicates sites that will be allowed, the right side are sites that will be blocked. In this case, only Social Networking sites with reputations in the Suspicious Sites and High Risk

ranges will be blocked. Thus, your users should be able to get to commonly-used Social Networking sites, where there are fewer risks.

Using reputation, you can selectively block sites within a category you otherwise want to allow.

- i) Click the + next to the **URLS** list to the left of the categories list.
- j) At the bottom of the popup dialog box, click the **Create New URL** link.
- k) Enter **badsite.example.com** for both the name and URL, then click **OK** to create the object.

You can name the object the same as the URL or give the object a different name. For the URL, do not include the protocol portion of the URL, just add the server name.

New URL Object

Name

Description

URL

- l) Select the new object, then click **OK**.

Adding new objects while editing policies simply adds the object to the list. The new object is not automatically selected.

Order	Title	Action
1	Block_Web_Sites	Block

Source/Destination Applications **URLS** Users Intrusion Policy File policy Logging

URLS	CATEGORIES
<input type="text" value="badsite.example.com"/>	<ul style="list-style-type: none"> <input type="checkbox"/> Adult and Pornograp... Reputation: Risk Any <input type="checkbox"/> Bot Nets Reputation: Risk Any <input type="checkbox"/> Confirmed SPAM So... Reputation: Risk Any <input type="checkbox"/> Social Network Reputation: Benign sites with sec...

- m) Click the **Logging** tab and select **Select Log Action > At Beginning and End of Connection**.

You must enable logging to get category and reputation information into the web category dashboard and connection events.

- n) Click **OK** to save the rule.

Step 3 (Optional.) Set preferences for URL filtering.

When you enable the URL license, the system automatically enables updates to the web category database. The system checks for updates every 30 minutes, although the data is typically updated once per day. You can turn off these updates if for some reason you do not want them.

You can also elect to send URLs that are not categorized to Cisco for analysis. Thus, if the installed URL database does not have a categorization for a site, the Cisco Cloud might have one. The cloud returns the category and reputation, and your category-based rules can then be applied correctly to the URL request. Selecting this option is important for lower-end systems, which install a smaller URL database due to memory limitations.

- a) Click **Device**.
- b) Click **System Settings > Traffic Settings > URL Filtering Preferences**.
- c) Select **Query Cisco CSI for Unknown URLs**.
- d) Click **Save**.

Step 4 Commit your changes.

- a) Click the **Deploy Changes** icon in the upper right of the web page.



- b) Click the **Deploy Now** button.

Wait for deployment to finish. The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.

What to do next

At this point, the monitoring dashboards and events should start showing information about web categories and reputations, and which connections were dropped. You can evaluate this information to determine if your URL filtering is dropping just those sites that are objectionable, or if you need to ease up on the reputation setting for certain categories.

Consider informing users beforehand that you will be blocking access to web sites based on their categorization and reputation.

How to Control Application Usage

The Web has become the ubiquitous platform for application delivery in the enterprise, whether that is browser based application platforms, or rich media applications that use web protocols as the transport in and out of enterprise networks.

FTD inspects connections to determine the application being used. This makes it possible to write access control rules targeted at applications, rather than just targeting specific TCP/UDP ports. Thus, you can selectively block or allow web-based applications even though they use the same port.

Although you can select specific applications to allow or block, you can also write rules based on type, category, tag, risk, or business relevance. For example, you could create an access control rule that identifies and blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the session is blocked.

Cisco frequently updates and adds additional application detectors via system and vulnerability database (VDB) updates. Thus, a rule blocking high risk applications can automatically apply to new applications without you having to update the rule manually.

In this use case, we will block any application that belongs to the **anonymizer/proxy** category.

Before you begin

This use case assumes that you completed the use case [How to Gain Insight Into Your Network Traffic, on page 36](#). That use case explains how to collect application usage information, which you can analyze in the Applications dashboard. Understanding what applications are actually being used can help you design effective application-based rules. The use case also explains how to schedule VDB updates, which will not be repeated here. Ensure that you update the VDB regularly so that applications can be correctly identified.

Procedure

Step 1

Create the application-based access control rule.

- a) Click **Policies** in the main menu.

Ensure that the **Access Control** policy is displayed.


- b) Click + to add a new rule.

- c) Configure the order, title, and action.


- **Order**—The default is to add new rules to the end of the access control policy. However, you must place this rule ahead of (above) any rule that would match the same Source/Destination and other criteria, or the rule will never be matched (a connection matches one rule only, and that is the first rule it matches in the table). For this rule, we will use the same Source/Destination as the `Inside_Outside_Rule` created during initial device configuration. You might have created other rules as well. To maximize access control efficiency, it is best to have specific rules early, to ensure the quickest decision on whether a connection is allowed or dropped. For the purposes of this example, select **1** as the rule order.
- **Title**—Give the rule a meaningful name, such as `Block_Anonymizers`.
- **Action**—Select **Block**.

Order	Title	Action
1	Block_Anonymizers	Block

- d) On the **Source/Destination** tab, click + for **Source > Zones**, select `inside_zone`, then click **OK** in the zones dialog box.



- e) Using the same technique, select **outside_zone** for **Destination > Zones**.



- f) Click the **Applications** tab.
 g) Click the + for **Applications**, and then click the **Advanced Filter** link at the bottom of the popup dialog box.

Although you can create application filter objects beforehand and select them on the Application Filters list here, you can also specify criteria directly in the access control rule, and optionally save the criteria as a filter object. Unless you are writing a rule for a single application, it is easier to use the Advanced Filter dialog box to find applications and construct appropriate criteria.

As you select criteria, the Applications list at the bottom of the dialog box updates to show exactly which applications match the criteria. The rule you are writing applies to these applications.

Look at this list carefully. For example, you might be tempted to block all very high risk applications. However, as of this writing, TFPT is classified as very high risk. Most organizations would not want to block this application. Take the time to experiment with various filter criteria to see which applications match your selections. Keep in mind that these lists can change with every VDB update.

For purposes of this example, select anonymizers/proxies from the Categories list.

Filter Applications

? RESET FILTER

Filter Applications interface showing filter options:

- Risks:** Any
- Business Relevance:** Any
- Types:** Any
- Categories:** 1 selected
 - Search Categories
 - anonymizer/proxy (selected)
 - mobile application
 - VoIP
 - web services provider
 - e-commerce
- Tags:** Any selected
 - Search Tags
 - displays ads
 - not work related
 - high bandwidth
 - file sharing/transfer
 - share media

Filter the list of applications

33 Applications

Application	Description
All applications that match the filters (33)	
ASProxy	ASProxy open-source web proxy
After School	Anonymous messaging app.
Avocent	Registered with IANA on port 1078 tcp/udp.
Avoidr	Web based proxy compatible with many popular social networking sites.

h) Click **Add** in the Advanced Filters dialog box.

The filter is added and shown on the Applications tab.

Source/Destination Applications URLs Users Intrusion Policy

APPLICATIONS SAVE AS FILTER +

Categories: anonymizer/proxy

i) Click the **Logging** tab and select **Select Log Action > At Beginning and End of Connection**.

You must enable logging to get information about any connections blocked by this rule.

j) Click **OK** to save the rule.

Step 2

Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

Wait for deployment to finish. The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.

Step 3 Click **Monitoring** and evaluate the results.

You might now see dropped connections on the Applications widget on the **Network Overview** dashboard. Use the **All/Denied/Allowed** drop-down options to focus just on dropped applications.

The **Applications** dashboard shows these results as well. If someone tries to use these applications, you should be able to correlate the application with the user attempting the connection, assuming that you enable identity policies and require authentication.

How to Add a Subnet

If you have an available interface on your device, you can wire it to a switch (or another router) to provide services to another subnet.

There are many potential reasons you would add a subnet. For this use case, we will address the following typical scenario.

- The subnet is an inside network using the private network 192.168.2.0/24.
- The interface for the network has the static address 192.168.2.1. In this example, the physical interface is devoted to the network. Another option is to use an already-wired interface and create a subinterface for the new network.
- The device will provide addresses to workstations on the network using DHCP, using 192.168.2.2-192.168.2.254 as the address pool.
- Network access to other inside networks, and to the outside network, will be allowed. Traffic going to the outside network will use NAT to obtain a public address.




Note This example assumes the unused interface is not part of a bridge group. If it is currently a bridge group member, you must first remove it from the bridge group before following this procedure.

Before you begin

Physically connect the network cable to the interface and to the switch for the new subnet.

Procedure

Step 1 Configure the interface.

- a) Click **Device**, then click the link in the **Interfaces** summary.
- b) Hover over the **Actions** cell on the right side of the row for the interface you wired, and click the edit icon ()
- c) Configure the basic interface properties.

- **Name**—A unique name for the interface. For this example, **inside_2**.
- **Status**—Click the status toggle to enable the interface.
- **IPv4 Address** tab—Select **Static** for **Type**, then enter **192.168.2.1/24**.

Edit Physical Interface

Interface Name: Status:

Description:

IPv4 Address | IPv6 Address | Advanced Options

Type: IP Address and Subnet Mask: /

- d) Click **Save**.

The interface list shows the updated interface status and the configured IP address.

GigabitEthernet1/3	inside_2	<input checked="" type="checkbox"/>	192.168.2.1	STATIC
--------------------	----------	-------------------------------------	-------------	--------

Step 2 Configure the DHCP server for the interface.

- Click **Device**.
- Click **System Settings > DHCP Server**.
- Click the **DHCP Servers** tab.

The table lists any existing DHCP servers. If you are using the default configuration, the list includes one for the inside interface.

- Click + above the table.
- Configure the server properties.
 - **Enable DHCP Server**—Click this toggle to enable the server.
 - **Interface**—Select the interface on which you are providing DHCP services. In this example, select **inside_2**.
 - **Address Pool**—The addresses the server can supply to devices on the network. Enter 192.168.2.2-192.168.2.254. Make sure you do not include the network address (.0), the interface address (.1), or the broadcast address (.255). Also, if you need static addresses for any devices on the network, exclude those addresses from the pool. The pool must be a single continuous series of addresses, so choose static addresses from the beginning or ending of the range.

f) Click **Add**.

#	INTERFACE	ENABLED DHCP SERVER	ADDRESS POOL
1	inside	Enabled	192.168.1.5-192.168.1.254
2	inside_2	Enabled	192.168.2.2-192.168.2.254

Step 3 Add the interface to the inside security zone.

To write policies on an interface, the interface must belong to a security zone. You write policies for the security zones. Thus, as you add and remove interfaces in the zones, you automatically change the policies applied to the interface.

- Click **Objects** in the main menu.
- Select **Security Zones** from the objects table of contents.
- Hover over the **Actions** cell on the right side of the row for the **inside_zone** object, and click the edit icon (✎).
- Click + under **Interfaces**, select the **inside_2** interface, and click **OK** in the interfaces list.

e) Click **Save**.

#	NAME	INTERFACES
1	inside_zone	inside, inside_2
2	outside_zone	outside

Step 4 Create an access control rule that allows traffic between the inside networks.

Traffic is not automatically allowed between any interfaces. You must create access control rules to allow the traffic that you want. The only exception is if you allow traffic in the access control rule's default action. For the purposes of this example, we will assume you retained the block default action that the device setup wizard configures. Thus, you need to create a rule that will allow traffic between the inside interfaces. If you have already created a rule like this, skip this step.

- a) Click **Policies** in the main menu.

Ensure that the **Access Control** policy is displayed.

- b) Click + to add a new rule.

- c) Configure the order, title, and action.


- **Order**—The default is to add new rules to the end of the access control policy. However, you must place this rule ahead of (above) any rule that would match the same Source/Destination and other criteria, or the rule will never be matched (a connection matches one rule only, and that is the first rule it matches in the table). For this rule, we will use unique Source/Destination criteria, so adding the rule to the end of the list is acceptable.

- **Title**—Give the rule a meaningful name, such as Allow_Inside_Inside.

- **Action**—Select **Allow**.

Order	Title	Action
4	Allow_Inside_Inside	Allow

- d) On the **Source/Destination** tab, click + for **Source > Zones**, select **inside_zone**, then click **OK** in the zones dialog box.



- e) Using the same technique, select **inside_zone** for **Destination > Zones**.

A security zone must contain at least two interfaces to select the same zone for source and destination.

Source/Destination	Applications	URLs	Users	Intrusion Policy	File policy	Logging
SOURCE				DESTINATION		
Zones	Networks	Ports	Zones			
inside_zone	ANY	ANY	inside_zone			

- f) (Optional.) Configure intrusion and malware inspection.

Although the inside interfaces are in a trusted zone, it is typical for users to connect laptops to the network. Thus, a user might unknowingly bring a threat inside your network from an outside network or a Wi-Fi hot spot. Thus, you might want to scan for intrusions and malware in traffic that goes between your inside networks.

Consider doing the following.

- Click the **Intrusion Policy** tab, enable the intrusion policy, and use the slider to select the Balanced Security and Connectivity policy.
- Click the **File Policy** tab, then select the Block Malware All policy.

g) Click the **Logging** tab and select **Select Log Action > At Beginning and End of Connection**.

You must enable logging to get information about any connections that match this rule. Logging adds statistics to the dashboard as well as showing events in the event viewer.

h) Click **OK** to save the rule.

Step 5 Verify that required policies are defined for the new subnet.

By adding the interface to the inside_zone security zone, any existing policies for inside_zone automatically apply to the new subnet. However, take the time to inspect your policies and ensure that no additional policies are needed.

If you completed the initial device configuration, the following policies should already apply.

- **Access Control**—The Inside_Outside_Rule should allow all traffic between the new subnet and the outside network. If you followed the previous use cases, the policy also provides intrusion and malware inspection. You must have a rule that allows some traffic between the new network and the outside network, or users cannot access the Internet or other external networks.
- **NAT**—The InsideOutsideNATrule applies to any interface going to the outside interface, and applies interface PAT. If you kept this rule, traffic from the new network going to the outside will have the IP address translated to a unique port on the outside interface's IP address. If you do not have a rule that applies to all interfaces, or the inside_zone interfaces, when going to the outside interface, you might need to create one now.
- **Identity**—There is no default identity policy. However, if you followed previous use cases, you might have an identity policy that already requires authentication for the new network. If you do not have an identity policy that applies, create one now if you want to have user-based information for the new network.

Step 6 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

Wait for deployment to finish. The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.

What to do next

Verify that workstations on the new subnet are getting IP addresses using DHCP, and that they can reach other inside networks and the outside network. Use the monitoring dashboards and the event viewer to evaluate network usage.

More Examples

In addition to the examples in the Use Case chapter, there are example configurations in some of the chapters that explain specific services. You might find the following examples of interest.

Network Address Translation (NAT)**NAT for IPv4 addresses**

- [Providing Access to an Inside Web Server \(Static Auto NAT\), on page 213](#)
- [Single Address for FTP, HTTP, and SMTP \(Static Auto NAT-with-Port-Translation\), on page 215](#)
- [Different Translation Depending on the Destination \(Dynamic Manual PAT\), on page 221](#)
- [Different Translation Depending on the Destination Address and Port \(Dynamic Manual PAT\), on page 227](#)
- [DNS Reply Modification, DNS Server on Outside, on page 239](#)
- [DNS Reply Modification, DNS Server on Host Network, on page 242](#)
- [Exempting Site-to-Site VPN Traffic from NAT, on page 266](#)

NAT for IPv6 addresses

- [NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet, on page 199](#)
- [NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet and DNS Translation, on page 201](#)
- [NAT66 Example, Static Translation between Networks, on page 206](#)
- [NAT66 Example, Simple IPv6 Interface PAT, on page 209](#)
- [DNS 64 Reply Modification, on page 233](#)

Site-to-Site Virtual Private Network (VPN)

- [Exempting Site-to-Site VPN Traffic from NAT, on page 266](#)
- [How to Provide Internet Access on the Outside Interface for External Site-to-Site VPN Users \(Hair Pinning\), on page 272](#)



CHAPTER 3

Licensing the System

The following topics explain how to license the FTD device.

- [Smart Licensing for the Firewall System, on page 63](#)
- [Managing Smart Licenses, on page 65](#)

Smart Licensing for the Firewall System

Cisco Smart Licensing is a flexible licensing model that provides you with an easier, faster, and more consistent way to purchase and manage software across the Cisco portfolio and across your organization. And it's secure—you control what users can access. With Smart Licensing you get:

- **Easy Activation:** Smart Licensing establishes a pool of software licenses that can be used across the entire organization—no more PAKs (Product Activation Keys).
- **Unified Management:** My Cisco Entitlements (MCE) provides a complete view into all of your Cisco products and services in an easy-to-use portal, so you always know what you have and what you are using.
- **License Flexibility:** Your software is not node-locked to your hardware, so you can easily use and transfer licenses as needed.

To use Smart Licensing, you must first set up a Smart Account on Cisco Software Central (software.cisco.com).

For a more detailed overview on Cisco Licensing, go to cisco.com/go/licensingguide

Cisco Smart Software Manager

When you purchase one or more licenses for the FTD device, you manage them in the Cisco Smart Software Manager: <https://software.cisco.com/#SmartLicensing-Inventory>. The Cisco Smart Software Manager lets you create a primary account for your organization.

By default, your licenses are assigned to the Default Virtual Account under your primary account. As the account administrator, you can create additional virtual accounts; for example, for regions, departments, or subsidiaries. Multiple virtual accounts help you manage large numbers of licenses and appliances.

Licenses and appliances are managed per virtual account; only that virtual account's appliances can use the licenses assigned to the account. If you need additional licenses, you can transfer an unused license from another virtual account. You can also transfer appliances between virtual accounts.

When you register a device with Cisco Smart Software Manager, you create a Product Instance Registration Token in the manager, and then enter it in FDM. A registered device becomes associated with a virtual account based on the token that is used.

For more information about the Cisco Smart Software Manager, see the online help for the manager.

Periodic Communication with the License Authority

When you use a Product Instance Registration Token to register the FTD device, the device registers with the Cisco License Authority. The License Authority issues an ID certificate for communication between the device and the License Authority. This certificate is valid for one year, although it will be renewed every six months. If an ID certificate expires (usually in nine months or a year with no communication), the device reverts to a de-registered state and licensed feature usage is suspended.

The device communicates with the License Authority on a periodic basis. If you make changes in the Cisco Smart Software Manager, you can refresh the authorization on the device so the changes immediately take effect. You also can wait for the device to communicate as scheduled. Normal license communication occurs every 12 hours, but with the grace period, your device will operate for up to 90 days without calling home. You must contact the License Authority before 90 days have passed.

Smart License Types

The following table explains the licenses available for the FTD device.

Your purchase of a FTD device automatically includes a Base license. All additional licenses are optional.

Table 2: Smart License Types

License	Duration	Granted Capabilities
Base	Perpetual	All features not covered by the optional term licenses. The Base license is automatically added to your account when you register. You must also specify whether to Allow export-controlled functionality on the products registered with this token . You can select this option only if your country meets export-control standards. This option controls your use of advanced encryption and the features that require advanced encryption.
Threat	Term-based	Required to use the following policies: <ul style="list-style-type: none"> • Intrusion • File (the Malware is also required)
Malware	Term-based	File policies (the Threat is also required).
URL	Term-based	Category and reputation-based URL filtering. You can perform URL filtering on individual URLs without this license.

Impact of Expired or Disabled Optional Licenses

If one of the following optional licenses expires, you can continue using features that require the license. However, the license is marked out of compliance and you need to purchase the license and add it to your account to bring the license back into compliance.

If you disable an optional license, the system reacts as follows:

- **Malware**—The system stops querying the Secure Malware Analytics Cloud, and also stops acknowledging retrospective events sent from the Secure Malware Analytics Cloud. You cannot re-deploy existing access control policies if they include file policies. Note that for a very brief time after a Malware license is disabled, the system can use existing cached file dispositions. After the time window expires, the system assigns a disposition of Unavailable to those files.
- **Threat**—The system no longer applies intrusion or file policies. You cannot re-deploy existing policies that require the license.
- **URL**—Access control rules with URL category conditions immediately stop filtering URLs, and the system no longer downloads updates to URL data. You cannot re-deploy existing access control policies if they include rules with category and reputation-based URL conditions.

Managing Smart Licenses

Use the Smart License page to view the current license status for the system. The system must be licensed.

The page shows you whether you are using the 90-day evaluation license, or if you have registered with the Cisco Smart Software Manager. Once registered, you can see the status of the connection to the Cisco Smart Software Manager as well as the status for each type of license.

Usage Authorization identifies the Smart License Agent status:

- **Authorized** (“Connected,” “Sufficient Licenses”)—The device has contacted and registered successfully with the License Authority, which has authorized the license entitlements for the appliance. The device is now In-Compliance.
- **Out-of-Compliance**—There is no available license entitlement for the device. Licensed features continue to work. However, you must either purchase or free up additional entitlements to become In-Compliance.
- **Authorization Expired**—The device has not communicated with the Licensing Authority in 90 or more days. Licensed features continue to work. In this state, the Smart License Agent retries its authorization requests. If a retry succeeds, the agent enters either an Out-of-Compliance or Authorized state, and begins a new Authorization Period. Try manually synchronizing the device.



Note Click the **i** button next to the Smart License status to view the virtual account, export-controlled features, and get a link to open the Cisco Smart Software Manager. Export-Controlled Features control software that is subject to national security, foreign policy, and anti-terrorism laws and regulations.

The following procedure provides an overview of how to manage licenses for the system.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Smart License summary.

Step 2 Register the device.

You must register with the Cisco Smart Software Manager before you can assign the optional licenses. Register before the end of the evaluation period.

See [Registering the Device, on page 66](#).

Step 3 Request and manage the optional feature licenses.

You must register the optional licenses to use the features controlled by the license. See [Enabling or Disabling Optional Licenses, on page 67](#).

Step 4 Maintain system licensing.

You can do the following tasks:

- [Synchronizing with the Cisco Smart Software Manager, on page 67](#)
 - [Unregistering the Device, on page 68](#)
-

Registering the Device

Your purchase of the FTD device automatically includes the Base license. The Base license covers all features not covered by the optional licenses. It is a perpetual license.

During initial system setup, you are prompted to register the device with Cisco Smart Software Manager. If you instead elected to use the 90-day evaluation license, you must register the device before the end of the evaluation period.

When you register the device, your virtual account allocates the license to the device. Registering the device also registers any optional licenses that you have enabled.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Smart License summary.

Step 2 Click **Request Register** and follow the instructions.

- a) Click the link to open the [Cisco Smart Software Manager](#) and log into your account, or create a new one if necessary.
- b) Generate a new token.

When you create the token, you specify the amount of time the token is valid for use. The recommended expiration period is 30 days. This period defines the expiration date of the token itself, and has no impact on the device that you register using the token. If the token expires before you can use it, you can simply generate a new token.

You must also specify whether to **Allow export-controlled functionality on the products registered with this token**. You can select this option only if your country meets export-control standards. This option controls your use of advanced encryption and the features that require advanced encryption.

- c) Copy and paste the token into the edit box on the Smart License Registration dialog box.
 - d) Click **Request Register**.
-

Enabling or Disabling Optional Licenses

You can enable (register) or disable (release) optional licenses. You must enable a license to use the features controlled by the license.

If you no longer want to use the features covered by an optional term license, you can disable the license. Disabling the license releases it in your Cisco Smart Software Manager account, so that you can apply it to another device.

You can also enable evaluation versions of these licenses when running in evaluation mode. In evaluation mode, the licenses are not registered with Cisco Smart Software Manager until you register the device.

Before you begin

Before disabling a license, ensure that you are not using it. Rewrite or delete any policies that require the license.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Smart License summary.

Step 2 Click the **Enable/Disable** control for each optional license as desired.

- **Enable**—Registers the license with your Cisco Smart Software Manager account and enables the controlled features. You can now configure and deploy policies controlled by the license.
 - **Disable**—Unregisters the license with your Cisco Smart Software Manager account and disables the controlled features. You cannot configure the features in new policies, nor can you deploy policies that use the feature.
-

Synchronizing with the Cisco Smart Software Manager

The system periodically synchronizes license information with Cisco Smart Software Manager. Normal license communication occurs every 30 days, but with the grace period, your appliance will operate for up to 90 days without calling home.

However, if you make changes in the Cisco Smart Software Manager, you can refresh the authorization on the device so the changes immediately take effect.

Synchronization gets the current status of licenses, and renews authorization and the ID certificate.

Procedure

- Step 1** Click **Device**, then click **View Configuration** in the Smart License summary.
- Step 2** Select **Resync Connection** from the gear drop-down list.
-

Unregistering the Device

If you no longer want to use the device, you can unregister it from the Cisco Smart Software Manager. When you unregister, the Base license and all optional licenses associated with the device are freed in your virtual account. Optional licenses are available to be assigned to other devices.

After unregistering the device, the current configuration and policies on the device continue to work as-is, but you cannot make or deploy any changes.

Procedure

- Step 1** Click **Device**, then click **View Configuration** in the Smart License summary.
- Step 2** Select **Unregister Device** from the gear drop-down list.
- Step 3** Read the warning and click **Unregister** if you really want to unregister the device.
-



CHAPTER 4

Monitoring the Device

The system includes dashboards and an Event Viewer that you can use to monitor the device and traffic that is passing through the device.

- [Enable Logging to Obtain Traffic Statistics, on page 69](#)
- [Monitoring Traffic and System Dashboards, on page 71](#)
- [Monitoring Additional Statistics Using the Command Line, on page 73](#)
- [Viewing Events, on page 74](#)

Enable Logging to Obtain Traffic Statistics

You can monitor a wide range of traffic statistics using the monitoring dashboards and the Event Viewer. However, you must enable logging to tell the system which statistics to collect. Logging generates various types of events that provide insight into the connections going through the system.

The following topics explain more about events and the information they provide, with special emphasis on connection logging.

Event Types

The system can generate the following types of events. You must generate these events to see related statistics in the monitoring dashboards.

Connection Events

You can generate events for connections as users generate traffic that passes through the system. Enable connection logging on access rules to generate these events.

Connection events include a wide variety of information about a connection, including source and destination IP addresses and ports, URLs and applications used, and the number of bytes or packets transmitted. The information also includes the action taken (for example, allowing or blocking the connection), and the policies applied to the connection.

Intrusion Events

The system examines the packets that traverse your network for malicious activity that could affect the availability, integrity, and confidentiality of a host and its data. When the system identifies a possible intrusion, it generates an intrusion event, which is a record of the date, time, type of exploit, and contextual information about the source of the attack and its target. Intrusion events are generated for any intrusion rule set to block or alert, regardless of the logging configuration of the invoking access control rule.

File Events

File events represent files that the system detected, and optionally blocked, in network traffic based on your file policies. You must enable file logging on the access rule that applies the file policy to generate these events.

When the system generates a file event, the system also logs the end of the associated connection regardless of the logging configuration of the invoking access control rule.

Malware Events

The system can detect malware in network traffic as part of your overall access control configuration. The AMP for Networks can generate a malware event, containing the disposition of the resulting event, and contextual data about how, where, and when the malware was detected. You must enable file logging on the access rule that applies the file policy to generate these events.

The disposition of a file can change, for example, from clean to malware or from malware to clean. If AMP for Networks queries the AMP Cloud about a file, and the cloud determines the disposition has changed within a week of the query, the system generates retrospective malware events.

Configurable Connection Logging

You should log connections according to the security and compliance needs of your organization. If your goal is to limit the number of events you generate and improve performance, only enable logging for the connections critical to your analysis. However, if you want a broad view of your network traffic for profiling purposes, you can enable logging for additional connections.

Because the system can log a connection for multiple reasons, disabling logging in one place does not ensure that matching connections will not be logged.

You configure connection logging on access control rules and the default action. Logging at the end of a connection provides the most information about the connection. You can also log the beginning of the connection, but these events have incomplete information. Connection logging is disabled by default, so you must enable it for each rule (and the default action) that targets traffic that you want to track.

Automatic Connection Logging

The system automatically saves the following end-of-connection events, regardless of any other logging configurations.

- The system automatically logs connections associated with intrusion events, unless the connection is handled by the access control policy's default action. You must enable logging on the default action to get intrusion events for matching traffic.
- The system automatically logs connections associated with file and malware events. This is for connection events only: you can optionally disable the generation of file and malware events.

Tips for Connection Logging

Keep the following tips in mind when considering your logging configuration and the evaluation of related statistics:

- When you allow traffic with an access control rule, you can use an associated intrusion or file policy (or both) to further inspect traffic and block intrusions, prohibited files, and malware before the traffic can

reach its final destination. Note, however, that by default file and intrusion inspection is disabled for encrypted payloads. If the intrusion or file policies find reason to block a connection, the system immediately logs an end-of-connection event regardless of your connection log settings. Logging allowed connections provides the most statistical information on the traffic in your network.

- A trusted connection is one that is handled by a Trust access control rule or the default action in an access control policy. However, trusted connections are not inspected for discovery data, intrusions, or prohibited files and malware. Therefore, connection events for trusted connections contain limited information.
- For access control rules and access control policy default actions that block traffic, the system logs beginning-of-connection events. Matching traffic is denied without further inspection.
- Logging blocked TCP connections during a Denial of Service (DoS) attack can affect system performance and overwhelm the database with multiple similar events. Before you enable logging for a Block rule, consider whether the rule monitors traffic on an Internet-facing interface or other interface vulnerable to DoS attack.

Sending Events to an External Syslog Server

Besides viewing events through the FDM, which has a limited capacity to store events, you can selectively configure rules and policies to send events to an external syslog server. You can then use the features and additional storage of your selected syslog server platform to view and analyze event data.

To send events to an external syslog server, edit each rule, default action, or policy that enables connection logging and select a syslog server object in the log settings.

For more information, see the help for each rule and policy type and also see [Configuring Syslog Servers, on page 94](#).

Monitoring Traffic and System Dashboards

The system includes several dashboards that you can use to analyze the traffic going through the device and the results of your security policy. Use the information to evaluate the overall efficacy of your configuration and to identify and resolve network problems.



Note The data used in traffic-related dashboards is collected from access control rules that enable connection or file logging. The dashboards do not reflect traffic that matches rules for which no logging is enabled. Ensure that you configure your rules to log the information that matters to you. In addition, user information is available only if you configure identity rules to collect user identity. And finally, intrusion, file, malware, and URL category information is available only if you have a license for those features and configure rules that use the features.

Procedure

Step 1 Click **Monitoring** in the main menu to open the Dashboards page.

You can select predefined time ranges, such as the last hour or week, or define a custom time range with specific start and end times, to control the data shown in the dashboard graphs and tables.

Traffic-related dashboards include the following types of display:

- **Top 5 bar graphs**—These are shown in the **Network Overview** dashboard, and in the per-item summary dashboards you see if you click on an item in a dashboard table. You can toggle the information between a count of **Transactions** or **Data Usage** (total bytes sent and received). You can also toggle the display to show all transactions, allowed transactions, or denied transactions. Click the **View More** link to see the table associated with the graph.
- **Tables**—Tables show items of a particular type (for example, applications or web categories) with that item's total transactions, allowed transactions, blocked transactions, data usage, and bytes sent and received. You can toggle the numbers between raw **Values** and **Percentages**, and show the top 10, 100, or 1000 entries. If the item is a link, click it to see a summary dashboard with more detailed information.

Step 2 Click the **Dashboard** links in the table of contents to see dashboards for the following data:

- **Network Overview**—Shows summary information about the traffic in the network, including the access rules (policies) matched, users initiating traffic, applications used in connections, intrusion threats (signatures) matched, web categories for URLs accessed, and the most frequent destinations for connections.
- **Users**—Shows the top users of your network. You must configure identity policies to see user information. You might see the following special entities:
 - **Failed Authentication**—The user was prompted to authenticate, but failed to enter a valid username/password pair within the maximum number of allowed attempts. Failure to authenticate does not itself prevent the user from accessing the network, but you can write an access rule to limit network access for these users.
 - **Guest**—Guest users are like Failed Authentication users, except that your identity rule is configured to call these users Guest. Guest users were prompted to authenticate and failed to do so within the maximum number of attempts.
 - **No Authentication Required**—The user was not prompted to authentication, because the user's connections matched identity rules that specified no authentication.
 - **Unknown**—There is no user mapping for the IP address, and there is no record of failed authentication yet. Typically, this means that no HTTP traffic has yet been seen from that address.
- **Applications**—Shows the top applications, such as HTTP, that are being used in the network. The information is available only for connections that are inspected. Connections are inspected if they match an “allow” rule, or a block rule that uses criteria other than zone, address, and port. Thus, application information is not available if the connection is trusted or blocked prior to hitting any rule that requires inspection.
- **Web Categories**—Shows the top categories of web sites, such as Gambling or Educational Institutions, that are being used in the network based on the categorization of web sites visited. You must have at least one access control rule that uses URL category as a traffic matching criteria to get this information. The information will be available for traffic that matches the rule, or for traffic that has to be inspected to determine if it matches the rule. You will not see category (or reputation) information for connections that match rules that come before the first web-category access control rule.
- **Policies**—Shows the top access rules matched by network traffic.
- **Ingress Zones**—Shows the top security zones through which traffic is entering the device.
- **Egress Zones**—Shows the top security zones through which traffic is exiting the device.

- **Destinations**—Shows the top destinations for network traffic.
- **Attackers**—Shows the top attackers, which are the source of connections that trigger intrusion events. You must configure intrusion policies on access rules to see this information.
- **Targets**—Shows the top targets of intrusion events, which are the victims of an attack. You must configure intrusion policies on access rules to see this information.
- **Threats**—Shows the top intrusion rules that have been triggered. You must configure intrusion policies on access rules to see this information.
- **File Logs**—Shows the top file types seen in network traffic. You must configure file policies on access rules to see this information.
- **System**— Shows an overall system view, including a display of interfaces and their status (mouse over an interface to see its IP addresses), overall average system throughput (in 5 minute buckets for up to one hour, and one hour buckets for longer periods), and summary information on system events, CPU usage, memory usage, and disk usage. You can restrict the throughput graph to show a specific interface rather than all interfaces. Interface-related statistics such as throughput does not include subinterfaces.

Note The information shown on the System dashboard is at the overall system level. If you log into the device CLI, you can use various commands to see more detailed information. For example, the **show cpu** and **show memory** commands include parameters for showing other details, whereas these dashboards show data from the **show cpu system** and **show memory system** commands.

Step 3 You can also click these links in the table of contents:

- **Events**—To view events as they occur. You must enable connection logging in individual access rules to see connection events related to those rules. These events can help you resolve connection problems for your users.

Monitoring Additional Statistics Using the Command Line

The FDM dashboards provide a wide variety of statistics related to the traffic going through the device and general system usage. However, you can get additional information on areas not covered by the dashboards by logging into the device CLI (see [Logging Into the Command Line Interface \(CLI\)](#), on page 5).

The CLI includes a variety of **show** commands to provide these statistics. You can also use the CLI for general troubleshooting, including commands such as **ping** and **traceroute**. Most **show** commands have companion **clear** commands to reset statistics to 0.

You can find documentation for the commands in [Cisco Firepower Threat Defense Command Reference](http://www.cisco.com/c/en/us/td/docs/security/firepower/command_ref/b_Command_Reference_for_Firepower_Threat_Defense.html), http://www.cisco.com/c/en/us/td/docs/security/firepower/command_ref/b_Command_Reference_for_Firepower_Threat_Defense.html.

For example, you might find the following commands generally useful.

- **show nat** displays hit counts for your NAT rules.
- **show xlate** displays the actual NAT translations that are active.
- **show conn** provides information about current connections going through the device.

- **show dhcpd** provides information about the DHCP servers you configure on the interfaces.
- **show interface** provides usage statistics for each interface.

Viewing Events

You can view events that are generated from your security policies that enable logging. Events are also generated for intrusion and file policies that are triggered.

The event viewer table shows the events generated in real time. As new events are generated, older events are rolled out of the table.

Before you begin

Whether events of particular types are generated depends on the following in addition to connections that match the related policy:

- Connection events—An access rule must enable connection logging.
- Intrusion events—An access rule must apply an intrusion policy.
- File and Malware events—An access rule must apply a file policy and enable file logging.

Procedure

Step 1 Click **Monitoring** in the main menu.

Step 2 Select **Events** from the table of contents.

The event viewer organizes events on tabs based on event types. For more information, see [Event Types, on page 69](#).

Step 3 Click the tab that shows the type of event you want to view.

You can do the following with the event list:

- Click **Pause** to stop the addition of new events so that you can more easily find and analyze an event. Click **Resume** to allow new events to appear.
- Select a different refresh rate (5, 10, 20, or 60 seconds) to control how fast new events are shown.
- Create a custom view that includes the columns you want. To create a custom view, either click the + button in the tab bar, or click **Add/Remove Columns**. You cannot change the pre-set tabs, so adding or removing columns creates a new view. For more information, see [Configuring Custom Views, on page 75](#).
- To change the width of a column, click and drag the column heading divider to the desired width.
- Mouse over an event and click **View Details** to see complete information on an event. For a description of the various fields in an event, see [Event Field Descriptions, on page 77](#).

Step 4 If necessary, apply a filter to the table to help you locate the desired events based on various event attributes.

To create a new filter, either manually type in the filter by selecting atomic elements from the drop-down list and entering the filter value, or build a filter by clicking a cell in the events table that includes a value on which you want to filter. You can click multiple cells in the same column to create an OR condition among the values, or click cells in different columns to create an AND condition among the columns. If you build the filter by clicking cells, you can also edit the resulting filter to fine-tune it. For detailed information about creating filter rules, see [Filtering Events, on page 76](#).

Once you build the filter, do any of the following:

- To apply the filter and update the table to show only those events that match the filter, click the **Filter** button.
- To clear an entire filter that you have applied and return the table to a non-filtered state, click **Reset Filters** in the **Filter** box.
- To clear one of the atomic elements of a filter, mouse over the element and click the **X** for the element. Then, click the **Filter** button.

Configuring Custom Views

You can create your own custom views so that you can easily see the columns you want when viewing events. You can also edit or delete custom views, although you cannot edit or delete the pre-defined views.

Procedure

Step 1 Select **Monitoring > Events**.

Step 2 Do one of the following:

- To create a new view based on an existing custom (or pre-defined) view, click the tab for the view, then click the + button to the left of the tabs.
- To edit an existing custom view, click the tab for the view.

Note To delete a custom view, simply click the **X** button in the view's tab. You cannot undo a delete.

Step 3 Click the **Add/Remove Columns** link above the events table on the right, and select or deselect columns until the selected list includes only those columns to include in the view.

Click and drag columns between the available (but not used) and selected lists. You can also click and drag columns in the selected list to change the left-to-right order of the columns in the table. For a description of the columns, see [Event Field Descriptions, on page 77](#).

When finished, click **OK** to save your column changes.

Note If you change column selection while viewing a pre-defined view, a new view is created.

Step 4 If necessary, change column widths by clicking and dragging the column separators.

Filtering Events

You can create complex filters to limit the events table to the events that currently interest you. You can use the following techniques, alone or in combination, to build a filter:

Clicking columns

The easiest way to build a filter is to click on cells in the events table that contain the values on which you intend to filter. Clicking a cell updates the **Filter** field with a correctly-formulated rule for that value and field combination. However, using this technique requires that the existing list of events contains the desired values.

You cannot filter on all columns. If you can filter on the contents of a cell, it is underlined when you mouse over it.

Selecting atomic elements

You can also build a filter by clicking in the **Filter** field and selecting the desired atomic element from the drop-down list, then typing in the match value. These elements include event fields that are not shown as columns in the events table. They also include operators to define the relationship between the value you type in and the events to display. Whereas clicking columns always results in an “equals (=)” filter, when you select an element, you can also select “greater than (>)” or “less than (<)” for numeric fields.

Regardless of how you add an element to the **Filter** field, you can type into the field to adjust the operator or value. Click **Filter** to apply the filter to the table.

Operators for Event Filters

You can use the following operators in an event filter:

=	Equals. The event matches the specified value. You cannot use wildcards.
!=	Not equals. The event does not match the specified value. You must type in the ! (exclamation point) to build a not-equals expression.
>	Greater than. The event contains a value that is greater than the specified value. This operator is available for numeric values only, such as port and IP address.
<	Less than. The event contains a value that is less than the specified value. This operator is available for numeric values only.

Rules for Complex Event Filters

When building a complex filter that contains more than one atomic element, keep the following rules in mind:

- Elements of the same type have an OR relationship between all values for that type. For example, including Initiator IP=10.100.10.10 and Initiator IP=10.100.10.11 matches events that have either of these addresses as the traffic source.
- Elements of different types have an AND relationship. For example, including Initiator IP=10.100.10.10 and Destination Port/ICMP Type=80 matches events that have this source address AND destination port only. Events from 10.100.10.10 to a different destination port are not shown.
- Numeric elements, including IPv4 and IPv6 addresses, can specify ranges. For example, you could specify Destination Port=50-80 to capture all traffic for ports within this range. Use a hyphen to separate the start and end numbers. Ranges are not allowed for all numeric fields, for example, you cannot specify an IP address range in the Source element.

- You cannot use wildcards or regular expressions.

Event Field Descriptions

Events can contain the following information. You can see this information when you view event details. You can also add columns to the Event Viewer table to show the information that most interests you.

Following is a complete list of the available fields. Not every field applies to every type of event. Keep in mind that the information available for any individual event can vary depending on how, why, and when the system logged the connection.

Action

For connection events, the action associated with the access control rule or default action that logged the connection:

Allow

Explicitly allowed connections.

Trust

Trusted connections. TCP connections detected by a trust rule on the first packet only generate an end-of-connection event. The system generates the event one hour after the final session packet.

Block

Blocked connections. The **Block** action can be associated with Allow access rules under the following conditions:

- Connections where an exploit was blocked by an intrusion policy.
- Connections where a file was blocked by a file policy.

Default Action

The connection was handled by the default action.

For file or malware events, the file rule action associated with the rule action for the rule the file matched, and any associated file rule action options.

Allowed Connection

Whether the system allowed the traffic flow for the event.

Application

The application detected in the connection.

Application Business Relevance

The business relevance associated with the application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of application detected in the connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Application Categories, Application Tag

Criteria that characterize the application to help you understand the application's function.

Application Risk

The risk associated with the application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of application detected in the connection has an associated risk; this field displays the highest of those.

Block Type

The type of block specified in the access control rule matching the traffic flow in the event: block or interactive block.

Client Application, Client Version

The client application and version of that client detected in the connection.

Client Business Relevance

The business relevance associated with the client traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of client detected in the connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Client Category, Client Tag

Criteria that characterize the application to help you understand the application's function.

Client Risk

The risk associated with the client traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of client detected in the connection has an associated risk; this field displays the highest of those.

Connection

The unique ID for the traffic flow, internally generated.

Connection Blocktype Indicator

The type of block specified in the access control rule matching the traffic flow in the event: block or interactive block.

Connection Bytes

The total bytes for the connection.

Connection Time

The time for the beginning of the connection.

Connection Timestamp

The time the connection was detected.

Denied Connection

Whether the system denied the traffic flow for the event.

Destination Country and Continent

The country and continent of the receiving host.

Destination IP

The IP address used by the receiving host in an intrusion, file, or malware event.

Destination Port/ICMP Code; Destination Port; Destination Icode

The port or ICMP code used by the session responder.

Direction

The direction of transmission for a file.

Disposition

The file's disposition:

Malware

Indicates that the AMP Cloud categorized the file as malware or the file's threat score exceeded the malware threshold defined in the file policy. Local malware analysis can also mark files as malware.

Clean

Indicates that the AMP Cloud categorized the file as clean, or that a user added the file to the clean list.

Unknown

Indicates that the system queried the AMP Cloud, but the file has not been assigned a disposition; in other words, the AMP Cloud has not categorized the file.

Unavailable

Indicates that the system could not query the AMP Cloud. You may see a small percentage of events with this disposition; this is expected behavior.

N/A

Indicates that a Detect Files or Block Files rule handled the file and the system did not query the AMP Cloud.

Egress Interface, Egress Security Zone

The interface and zone through which the connection exited the device.

Event, Event Type

The type of event.

Event Seconds, Event Microseconds

The time, in seconds or microseconds, when the event was detected.

File Category

The general categories of file type, for example: Office Documents, Archive, Multimedia, Executables, PDF files, Encoded, Graphics, or System Files.

File Event Timestamp

The time and date the file or malware file was created.

File Name

The name of the file.

File Rule Action

The action associated with file policy rule that detected the file, and any associated file rule action options.

File SHA-256

The SHA-256 hash value of the file.

File Size (KB)

The size of the file, in kilobytes. File size can be blank in cases where the system blocked the file before it was completely received.

File Type

The type of file, for example, HTML or MSEXEXE.

File/Malware Policy

The file policy associated with the generation of the event.

Filelog Blocktype Indicator

The type of block specified in the file rule matching the traffic flow in the event: block or interactive block.

Firewall Policy Rule, Firewall Rule

The access control rule or default action that handled the connection.

First Packet

The date and time the first packet of the session was seen.

HTTP Referrer

The HTTP referrer, which represents the referrer of a requested URL for HTTP traffic detected in the connection (such as a website that provided a link to, or imported a link from, another URL).

HTTP Response

The HTTP status code sent in response to a client's HTTP request over a connection.

IDS Classification

The classification where the rule that generated the event belongs.

Ingress Interface, Ingress Security Zone

The interface and zone through which the connection entered the device.

Initiator Bytes, Initiator Packets

The total number of bytes or packets transmitted by the session initiator.

Initiator Country and Continent

The country and continent of the host that initiated the session. Available only if the initiator IP address is routable.

Initiator IP

The host IP address (and hostname, if DNS resolution is enabled) that initiated the session in a connection or Security Intelligence event.

Inline Result

Whether the system dropped or would have dropped the packet that triggered an intrusion event if operating in inline mode. Blank indicates that the triggered rule was not set to Drop and Generate Events

Intrusion Policy

The intrusion policy where the rule that generated the event was enabled.

IPS Blocktype Indicator

The action of the intrusion rule matching the traffic flow in the event.

Last Packet

The date and time the last packet of the session was seen.

MPLS Label

The Multiprotocol Label Switching label associated with the packet that triggered this intrusion event.

Malware Blocktype Indicator

The type of block specified in the file rule matching the traffic flow in the event: block or interactive block.

Message

For intrusion events, the explanatory text for the event. For malware or file events, any additional information associated with the malware event.

NetBIOS Domain

The NetBIOS domain used in the session.

Original Client Country and Continent

The country and continent of the original client host that initiated the session. Available only if the original client IP address is routable.

Original Client IP

The original IP address of the client that initiated an HTTP connection. This address is derived from the X-Forwarded-For (XFF) or True-Client-IP HTTP header fields or their equivalent.

Policy, Policy Revision

The access control policy, and its revision, that includes the access (firewall) rule associated with the event.

Priority

The event priority as determined by the Cisco Talos Intelligence Group (Talos): high, medium, or low.

Protocol

The transport protocol used in the connection.

Reason

The reason or reasons the connection was logged, in the situations explained in the following table. This field is otherwise empty.

Reason	Description
File Block	The connection contained a file or malware file that the system prevented from being transmitted. A reason of File Block is always paired with an action of Block.
File Monitor	The system detected a particular type of file in the connection.

Reason	Description
File Resume Allow	File transmission was originally blocked by a Block Files or Block Malware file rule. After a new access control policy allowing the file was deployed, the HTTP session automatically resumed.
File Resume Block	File transmission was originally allowed by a Detect Files or Malware Cloud Lookup file rule. After a new access control policy blocking the file was deployed, the HTTP session automatically stopped.
Intrusion Block	The system blocked or would have blocked an exploit (intrusion policy violation) detected in the connection. A reason of Intrusion Block is paired with an action of Block for blocked exploits and Allow for would-have-blocked exploits.
Intrusion Monitor	The system detected, but did not block, an exploit detected in the connection. This occurs when the state of the triggered intrusion rule is set to Generate Events.

Receive Times

The date and time the event was generated.

Referenced Host

If the protocol in the connection is HTTP or HTTPS, this field displays the hostname that the respective protocol was using.

Responder Bytes, Responder Packets

The total number of bytes or packets transmitted by the session responder.

Responder Country and Continent

The country and continent of the host that responded to the session. Available only if the responder IP address is routable.

Responder IP

The host IP address (and hostname, if DNS resolution is enabled) of the session responder in a connection or Security Intelligence event.

Signature

The signature ID for a file/malware event.

Source Country and Continent

The country and continent of the sending host. Available only if the source IP address is routable.

Source IP

The IP address used by the sending host in an intrusion, file, or malware event.

Source Port/ICMP Type; Source Port; Source Port Itype

The port or ICMP type used by the session initiator.

TCP Flags

The TCP flags detected in the connection.

Total Packets

The total number of packets transmitted in the connection, which is **Initiator Packets + Responder Packets**.

URL, URL Category, URL Reputation, URL Reputation Score

The URL requested by the monitored host during the session and its associated category, reputation, and reputation score, if available.

If the system identifies or blocks an SSL application, the requested URL is in encrypted traffic, so the system identifies the traffic based on an SSL certificate. For SSL applications, therefore, the URL indicates the common name contained in the certificate.

User

The user associated with the initiator IP address.

VLAN

The innermost VLAN ID associated with the packet that triggered the event.

Web App Business Relevance

The business relevance associated with the web application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of web application detected in the connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Web App Categories, Web App Tag

Criteria that characterize the web application to help you understand the web application's function.

Web App Risk

The risk associated with the web application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of web application detected in the connection has an associated risk; this field displays the highest of those.

Web Application

The web application, which represents the content or requested URL for HTTP traffic detected in the connection.

If the web application does not match the URL for the event, the traffic is probably referred traffic, such as advertisement traffic. If the system detects referred traffic, it stores the referring application (if available) and lists that application as the web application.



CHAPTER 5

Objects

Objects are reusable containers that define criteria that you want to use in policies or other settings. For example, network objects define host and subnet addresses.

Objects let you define criteria so that you can easily reuse the same criteria in different policies. When you update an object, all policies that use the object are automatically updated.

- [Object Types, on page 85](#)
- [Managing Objects, on page 86](#)

Object Types

You can create the following types of object. In most cases, if a policy or setting allows an object, you must use an object.

Object Type	Main Use	Description
Application Filter	Access control rules.	An application filter object defines the applications used in an IP connection, or a filter that defines applications by type, category, tag, risk, or business relevance. You can use these objects in policies to control traffic instead of using port specifications. See Configuring Application Filter Objects, on page 90 .
Geolocation	Security policies.	A geolocation object defines countries and continents that host the device that is the source or destination of traffic. You can use these objects in policies to control traffic instead of using IP addresses. See Configuring Geolocation Objects, on page 94 .
IKE Policy	VPN.	Internet Key Exchange (IKE) Policy objects define the IKE proposal used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and automatically establish IPsec security associations (SAs). There are separate objects for IKEv1 and IKEv2. See Configuring the Global IKE Policy, on page 256 .

Object Type	Main Use	Description
IPsec Proposal	VPN.	IPsec Proposal objects configure the IPsec proposal used during IKE Phase 2 negotiations. The IPsec proposal defines the combination of security protocols and algorithms that secure traffic in an IPsec tunnel. There are separate objects for IKEv1 and IKEv2. See Configuring IPsec Proposals, on page 261 .
Network	Security policies and a wide variety of device settings.	Network groups and network objects (collectively referred to as network objects) define the addresses of hosts or networks. See Configuring Network Objects and Groups, on page 87 .
Port	Security policies.	Port groups and port objects (collectively referred to as port objects) define the protocols, ports, or ICMP services for traffic. See Configuring Port Objects and Groups, on page 88 .
Security Zone	Security policies.	A security zone is a grouping of interfaces. Zones divide the network into segments to help you manage and classify traffic. See Configuring Security Zones, on page 89 .
Syslog Servers	Access control rules. Diagnostic logging.	A syslog server object identifies a server that can receive connection-oriented or diagnostic system log (syslog) messages. See Configuring Syslog Servers, on page 94 .
URL	Access control rules.	URL objects and groups (collectively referred to as URL objects) define the URL or IP addresses of web requests. See Configuring URL Objects and Groups, on page 92 .

Managing Objects

You can configure objects directly through the Objects page, or you can configure them while editing policies. Either method yields the same results, a new or updated object, so use the technique that suits your needs at the time.

The following procedure explains how you can create and manage your objects directly through the Objects page.






Note When you edit a policy or setting, if a property requires an object, you are shown a list of the ones that are already defined, and you select the appropriate object. If the desired object does not yet exist, simply click the **Create New Object** link shown in the list.

Procedure

Step 1 Select **Objects**.

The Objects page has a table of contents listing the available types of objects. When you select an object type, you see a list of existing objects, and you can create new ones from here. You can also see the object contents and type.

Step 2 Select the object type from the table of contents and do any of the following:

- To create an object, click the + button. The content of the objects differ based on type; see the configuration topic for each object type for specific information.
 - To create a group object, click the **Add Group** () button. Group objects include more than one item.
 - To edit an object, click the edit icon () for the object. You cannot edit the contents of a pre-defined object.
 - To delete an object, click the delete icon () for the object. You cannot delete an object if it is currently being used in a policy or another object, or if it is a pre-defined object.
-

Configuring Network Objects and Groups



Use network group and network objects (collectively referred to as network objects) to define the addresses of hosts or networks. You can then use the objects in security policies for purposes of defining traffic matching criteria, or in settings to define the addresses of servers or other resources.


A network object defines a single host or network address, whereas a network group object can define more than one address.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create network objects while editing an address property by clicking the **Create New Network** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **Network** from the table of contents.**Step 2** Do one of the following:

- To create an object, click the + button.
- To create a group, click the **Add Group** () button.
- To edit an object or group, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 3 Enter a Name for the object and optionally, a description, and define the object contents.

We recommend that you do not use an IP address alone for the name so that you can easily tell object names from object contents or standalone IP addresses. If you want to use an IP address in the name, prefix it with something meaningful, such as host-192.168.1.2 or network-192.168.1.0. If you use an IP address as the name,

the system adds a vertical bar as a prefix, for example, |192.168.1.2. FDM does not show the bar in the object selectors, but you will see this naming standard if you examine the running configuration using the **show running-config** command in the CLI.

Step 4 Configure the contents of the object.

Network Objects

Select the object **Type** and configure the contents:

- **Network**—Enter a network address using one of the following formats:
 - IPv4 network including subnet mask, for example, 10.100.10.0/24 or 10.100.10.0/255.255.255.0.
 - IPv6 network including prefix, for example, 2001:DB8:0:CD30::/60.
- **Host**—Enter a host IP address using one of the following formats:
 - IPv4 host address, for example, 10.100.10.10.
 - IPv6 host address, for example, 2001:DB8::0DB8:800:200C:417A or 2001:DB8:0:0:0DB8:800:200C:417A.

Network Groups

Click the + button to select network objects to add to the group. You can also create new objects.

Step 5 Click **OK** to save your changes.

Configuring Port Objects and Groups

Use port group and port objects (collectively referred to as port objects) to define the protocols, ports, or ICMP services for traffic. You can then use the objects in security policies for purposes of defining traffic matching criteria, for example, to use access rules to allow traffic to specific TCP ports.

A port object defines a single protocol, TCP/UDP port or port range, or ICMP service, whereas a port group object can define more than one service.

The system includes several pre-defined objects for common services. You can use these objects in your policies. However, you cannot edit or delete system-defined objects.



Note When creating port group objects, ensure that the combination of objects makes sense. For example, you cannot have a mixture of protocols in an object if you use it to specify both source and destination ports in an access rule. Exercise care when editing an object that is already being used, or you could invalidate (and disable) policies that use the object.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create port objects while editing a service property by clicking the **Create New Port** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **Ports** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To create a group, click the **Add Group** (📁) button.
- To edit an object or group, click the edit icon (✎) for the object.

To delete an unreferenced object, click the trash can icon (🗑️) for the object.

Step 3 Enter a name for the object and optionally, a description, and define the object contents.

Port Objects

Select the **Protocol**, then configure the protocol as follows:

- **TCP, UDP**—Enter the single port or port range number, for example, 80 (for HTTP) or 1-65535 (to cover all ports).
- **ICMP, IPv6-ICMP**—Select the ICMP **Type** and optionally, the **Code**. Select **Any** for the type to apply to all ICMP messages. For information on the types and codes, see the following pages:
 - ICMP—<http://www.iana.org/assignments/icmp-parameters/icmp-parameters.xml>
 - ICMPv6—<http://www.iana.org/assignments/icmpv6-parameters/icmpv6-parameters.xml>
- **Other**—Select the desired protocol.

Port Groups

Click the + button to select port objects to add to the group. You can also create new objects.

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

Configuring Security Zones

A security zone is a grouping of interfaces. Zones divide the network into segments to help you manage and classify traffic. You can define multiple zones, but a given interface can be in one zone only.

The system creates the following zones during initial configuration. You can edit these zones to add or remove interfaces, or you can delete the zones if you no longer use them.

- **inside_zone**—Includes the inside interface. If the inside interface is a bridge group, this zone includes all the bridge group member interfaces instead of the inside Bridge Virtual Interface (BVI). This zone is intended to represent internal networks.
- **outside_zone**—Includes the outside interface. This zone is intended to represent networks external to your control, such as the Internet.

Typically, you would group interfaces by the role they play in your network. For example, you would place the interface that connects to the Internet in the **outside_zone** security zone, and all of the interfaces for your

internal networks in the **inside_zone** security zone. Then, you could apply access control rules to traffic coming from the outside zone and going to the inside zone.


Before creating zones, consider the access rules and other policies you want to apply to your networks. For example, you do not need to put all internal interfaces into the same zone. If you have 4 internal networks, and you want to treat one differently than the other three, you can create two zones rather than one. If you have an interface that should allow outside access to a public web server, you might want to use a separate zone for the interface.


The following procedure explains how you can create and edit objects directly through the Objects page. You can also create security zones while editing a security zone property by clicking the **Create New Security Zone** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **Security Zones** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 3 Enter a Name for the object and optionally, a description.

Step 4 In the **Interfaces** list, click + and select the interfaces to add to the zone.

The list shows all named interfaces that are not currently in a zone. You must configure an interface and give it a name before you can add it to a zone.

If all named interfaces are already in zones, the list is empty. If you are trying to move an interface to a different zone, you must first remove it from its current zone.

Note You cannot add a bridge group interface (BVI) to a zone. Instead, add the member interfaces. You can put the members into different zones.

Step 5 Click **OK** to save your changes.

Configuring Application Filter Objects

An application filter object defines the applications used in an IP connection, or a filter that defines applications by type, category, tag, risk, or business relevance. You can use these objects in policies to control traffic instead of using port specifications.

Although you can specify individual applications, application filters simplify policy creation and administration. For example, you could create an access control rule that identifies and blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the session is blocked.

You can select applications and application filters directly in a policy without using application filter objects. However, an object is convenient if you want to create several policies for the same group of applications or filters. The system includes several pre-defined application filters, which you cannot edit or delete.



Note Cisco frequently updates and adds additional application detectors via system and vulnerability database (VDB) updates. Thus, a rule blocking high risk applications can automatically apply to new applications without you having to update the rule manually.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create application filter objects while editing an access control rule by clicking the **Save As Filter** link after adding application criteria to the Applications tab.


Before you begin


When editing a filter, if a selected application was removed by a VDB update, “(Deprecated)” appears after the application name. You must remove these applications from the filter, or subsequent deployments and system software upgrades will be blocked.

Procedure

Step 1 Select **Objects**, then select **Application Filters** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 3 Enter a Name for the object and optionally, a description.

Step 4 In the **Applications** list, click **Add +** and select the applications and filters to add to the object.

The initial list shows applications in a continually scrolling list. Click **Advanced Filter** to see the filter options and to get an easier view for selecting applications. Click **Add** when you have made your selections. You can repeat the process to add additional applications or filters.

Note Multiple selections within a single filter criteria have an OR relationship. For example, Risk is High OR Very High. The relationship between filters is AND, so Risk is High OR Very High, AND Business Relevance is Low OR Very Low. As you select filters, the list of applications in the display updates to show only those that meet the criteria. You can use these filters to help you find applications that you want to add individually, or to verify that you are selecting the desired filters to add to the rule.

Risks

The likelihood that the application is used for purposes that might be against your organization's security policy, from very low to very high.

Business Relevance

The likelihood that the application is used within the context of your organization's business operations, as opposed to recreationally, from very low to very high.

Types

The type of application:

- **Application Protocol**—Application protocols such as HTTP and SSH, which represent communications between hosts.
- **Client Protocol**—Clients such as web browsers and email clients, which represent software running on the host.
- **Web Application**—Web applications such as MPEG video and Facebook, which represent the content or requested URL for HTTP traffic.

Categories

A general classification for the application that describes its most essential function.

Tags

Additional information about the application, similar to category.

For encrypted traffic, the system can identify and filter traffic using only the applications tagged **SSL Protocol**. Applications without this tag can only be detected in unencrypted or decrypted traffic. Also, the system assigns the **decrypted traffic** tag to applications that the system can detect in decrypted traffic only, not encrypted or unencrypted.

Applications List (bottom of the display)

This list updates as you select filters from the options above the list, so you can see the applications that currently match the filter. Use this list to verify that your filter is targeting the desired applications when you intend to add filter criteria to the rule. If your intention is to add specific applications, select them from this list.

Step 5 Click **OK** to save your changes.

Configuring URL Objects and Groups

Use URL objects and groups (collectively referred to as URL objects) to define the URL or IP addresses of web requests. You can use these objects to implement manual URL filtering in access control policies.

A URL object defines a single URL or IP address, whereas a URL group object can define more than one URL or address.

When creating URL objects, keep the following points in mind:

- If you do not include a path (that is, there is no / character in the URL), the match is based on the server's hostname only. The hostname is considered a match if it comes after the :// separator, or after any dot in the hostname. For example, ign.com matches ign.com and www.ign.com, but it does not match verisign.com.
- If you include one or more / character, the entire URL string is used for a substring match, including the server name, path, and any query parameters. However, we recommend that you do not use manual URL filtering to block or allow individual web pages or parts of sites, as servers can be reorganized and pages moved to new paths. Substring matching can also lead to unexpected matches, where the string you include in the URL object also matches paths on unintended servers or strings within query parameters.

- The system disregards the encryption protocol (HTTP vs HTTPS). In other words, if you block a website, both HTTP and HTTPS traffic to that website is blocked, unless you use an application condition to target a specific protocol. When creating a URL object, you do not need to specify the protocol when creating an object. For example, use `example.com` rather than `http://example.com`.
- If you plan to use a URL object to match HTTPS traffic in an access control rule, create the object using the subject common name in the public key certificate used to encrypt the traffic. Also, the system disregards subdomains within the subject common name, so do not include subdomain information. For example, use `example.com` rather than `www.example.com`.

However, please understand that the subject common name in the certificate might be completely unrelated to a web site's domain name. For example, the subject common name in the certificate for `youtube.com` is `*.google.com` (this of course might change at any time). You will get more consistent results if you use the SSL Decryption policy to decrypt HTTPS traffic so that URL filtering rules work on decrypted traffic.





Note URL objects will not match HTTPS traffic if the browser resumes a TLS session because the certificate information is no longer available. Thus, even if you carefully configure the URL object, you might get inconsistent results for HTTPS connections.


The following procedure explains how you can create and edit objects directly through the Objects page. You can also create URL objects while editing a URL property by clicking the **Create New URL** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **URL** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To create a group, click the **Add Group** () button.
- To edit an object or group, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 3 Enter a Name for the object and optionally, a description.

Step 4 Define the object contents.

URL Objects

Enter a URL or IP address in the **URL** box. You cannot use wildcards in the URL.

URL Groups

Click the + button to select URL objects to add to the group. You can also create new objects.

Step 5 Click **OK** to save your changes.

Configuring Geolocation Objects

A geolocation object defines countries and continents that host the device that is the source or destination of traffic. You can use these objects in policies to control traffic instead of using IP addresses. For example, using geographical location, you could easily restrict access to a particular country without needing to know all of the potential IP addresses used there.

You can typically select geographical locations directly in a policy without using geolocation objects. However, an object is convenient if you want to create several policies for the same group of countries and continents.




Note To ensure that you are using up-to-date geographical location data to filter your traffic, Cisco strongly recommends that you regularly update the geolocation database (GeoDB).


The following procedure explains how you can create and edit objects directly through the Objects page. You can also create geolocation objects while editing a network property by clicking the **Create New Geolocation** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **Geolocation** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 3 Enter a Name for the object and optionally, a description.

Step 4 In the **Continents/Countries** list, click **Add +** and select the continents and countries to add to the object.

Selecting a continent selects all countries within the continent.

Step 5 Click **OK** to save your changes.

Configuring Syslog Servers

A syslog server object identifies a server that can receive connection-oriented or diagnostic system log (syslog) messages. If you have a syslog server set up for log collection and analysis, create objects to define them and use the objects in the related policies.

You can send the following types of events to the syslog server:


- Connection events. Configure the syslog server object on the following types of policy: access control rules and default action.
- Diagnostic events. See [Configuring Diagnostic Logging, on page 283](#).


The following procedure explains how you can create and edit objects directly through the Objects page. You can also create syslog server objects while editing a syslog server property by clicking the **Add Syslog Server** link shown in the object list.

Procedure

Step 1 Select **Objects**, then select **Syslog Servers** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 3 Configure the syslog server properties:

- **Device Interface**—Select the interface through which the syslog server is reached. If the server is accessible through a bridge group member interface, select the bridge group interface (BVI) instead.
- **IP Address**—Enter the IP address of the syslog server.
- **Port**—Enter the UDP port that the server uses for receiving syslog messages. The default is 514. If you change the default, the port must be in the range 1025 to 65535.

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



PART I

The Basics

- [Interfaces, on page 99](#)
- [Routing, on page 117](#)



CHAPTER 6

Interfaces

The following topics explain how to configure the interfaces on your FTD device.

- [About FTD Interfaces, on page 99](#)
- [Guidelines and Limitations for Interfaces, on page 102](#)
- [Configure a Physical Interface, on page 104](#)
- [Configure Bridge Groups, on page 107](#)
- [Configure VLAN Subinterfaces and 802.1Q Trunking, on page 110](#)
- [Configure Advanced Interface Options, on page 113](#)
- [Monitoring Interfaces, on page 116](#)
- [Examples for Interfaces, on page 116](#)

About FTD Interfaces

FTD includes data interfaces as well as a Management/Diagnostic interface.

When you attach a cable to an interface connection, you need to configure the interface. At minimum, you need to name the interface and enable it for it to pass traffic. If the interface is a member of a bridge group, this is sufficient. For non-bridge group members, you also need to give the interface an IP address. If you intend to create VLAN subinterfaces rather than a single physical interface on a given port, you would typically configure the IP addresses on the subinterface, not on the physical interface. VLAN subinterfaces let you divide a physical interface into multiple logical interfaces that are tagged with different VLAN IDs, which is useful when you connect to a trunk port on a switch.

The interface list shows the available interfaces, their names, addresses, and states. You can change the state of an interface, on or off, directly in the list of interfaces. The list shows the interface characteristics based on your configuration. Use the open/close arrow on a bridge group interface to view the member interfaces, which also appear by themselves in the list.

The following topics explain the limitations of configuring interfaces through the FDM as well as other interface management concepts.

Interface Modes

You can configure one of the following modes for each interface:

Routed

Each Layer 3 routed interface requires an IP address on a unique subnet. You would typically attach these interfaces to switches, a port on another router, or to an ISP/WAN gateway.

BridgeGroupMember

A bridge group is a group of interfaces that the FTD device bridges instead of routes. All interfaces are on the same network. The bridge group is represented by a Bridge Virtual Interface (BVI) that has an IP address on the bridge network.

You can route between routed interfaces and BVIs, if you name the BVI. In this case, the BVI acts as the gateway between member interfaces and routed interfaces. If you do not name the BVI, traffic on the bridge group member interfaces cannot leave the bridge group. Normally, you would name the interface so that you can route member interfaces to the internet.

One use for a bridge group in routed mode is to use extra interfaces on the FTD device instead of an external switch. You can attach endpoints directly to bridge group member interfaces. You can also attach switches to add more endpoints to the same network as the BVI.

Management/Diagnostic Interface

The physical port labeled Management actually has two separate interfaces associated with it.

- Management virtual interface—This IP address is used for system communication. This is the address the system uses for Smart Licensing and to retrieve database updates. You can open management sessions to it (FDM and CLI). You must configure a management address, which is defined on **System Settings > Management Interface**.
- Diagnostic virtual interface—You can use this interface to send syslog messages to an external syslog server. Configuring an IP address for the Diagnostic interface is optional. The main reason to configure the interface is if you want to use it for syslog messages. This interface appears, and is configurable, on the **Device > Interfaces** page. The Diagnostic interface only allows management traffic, and does not allow through traffic.

One way to configure Management/Diagnostic is to not wire the physical port to a network. Instead, configure the Management IP address only, and configure it to use the data interfaces as the gateway for obtaining updates from the internet. Then, open the inside interfaces to HTTPS/SSH traffic (by default, HTTPS is enabled) and open the FDM using the inside IP address (see [Configuring the Management Access List](#), on page 281).

Recommendations for Configuring a Separate Management Network

If you want to use a separate management network, wire the physical Management interface to a switch or router.

Then, configure the following:

- Select **Device > System Settings > Management Interface** and configure IPv4 or IPv6 addresses (or both) on the attached network. If you want to, you can configure a DHCP server to provide IPv4 addresses to other endpoints on the network. If there is a router with a route to the internet on the management network, use that as the gateway. Otherwise, use the data interfaces as the gateway.
- Configure an address for the Diagnostic interface (on **Device > Interfaces**) only if you intend to send syslog messages through the interface to a syslog server. Otherwise, do not configure an address for Diagnostic; it is not needed. Any IP address you configure must be on the same subnet as the management

IP address and cannot be the in DHCP server pool. For example, if you use 192.168.45.45 as the management address, and 192.168.45.46-192.168.45.254 as the DHCP pool, you can configure Diagnostic using any address from 192.168.45.1 to 192.168.45.44.

Limitations for Management/Diagnostic Interface Configuration for a Separate Management Network

If you wire the physical Management interface, ensure that you follow these limitations:

- If you want a DHCP server on the management network, configure it on the Management interface (**Device > System Settings > Management Interface**). You cannot configure a DHCP server on the Diagnostic interface.
- If there is another DHCP server on the management network, disable it or the DHCP server running on Management. As a rule, a given subnet should have no more than one DHCP server.
- If you configure addresses for both Management and Diagnostic, ensure that they are on the same subnet.
- You can use the data interfaces as the management gateway even if you configure an IP address for Diagnostic. But Diagnostic will not use the data interfaces as a gateway. If you need a path from Diagnostic to other networks, another router on the management network needs to route the traffic originating from the Diagnostic IP address. If necessary, configure static routes for the Diagnostic interface (select **Device > Routing**).

Security Zones

Each interface can be assigned to a single security zone. You then apply your security policy based on zones. For example, you can assign the inside interface to the inside zone; and the outside interface to the outside zone. You can configure your access control policy to enable traffic to go from inside to outside, but not from outside to inside, for example.

For bridge groups, you add member interfaces to the zones, you cannot add the Bridge Virtual Interface (BVI).

You do not include the Management/Diagnostic interface in a zone. Zones apply to data interfaces only.

You can create security zones on the **Objects** page.

IPv6 Addressing

You can configure two types of unicast addresses for IPv6:

- Global—The global address is a public address that you can use on the public network. For a bridge group, you configure the global address on the Bridge Virtual Interface (BVI), not on each member interface. You cannot specify any of the following as a global address.
 - Internally reserved IPv6 addresses: fd00::/56 (from=fd00:: to= fd00:0000:0000:00ff:ffff:ffff:ffff:ffff)
 - An unspecified address, such as ::/128
 - The loopback address, ::1/128
 - multicast addresses, ff00::/8
 - Link-local addresses, fe80::/10

- **Link-local**—The link-local address is a private address that you can only use on the directly-connected network. Routers do not forward packets using link-local addresses; they are only for communication on a particular physical network segment. They can be used for address configuration or for the Network Discovery functions such as address resolution and neighbor discovery. In a bridge group, enabling IPv6 on the BVI automatically configures link-local addresses for each bridge group member interface. Each interface must have its own address because the link-local address is only available on a segment, and is tied to the interface MAC address.

At a minimum, you need to configure a link-local address for IPv6 to operate. If you configure a global address, a link-local address is automatically configured on the interface, so you do not also need to specifically configure a link-local address. If you do not configure a global address, then you need to configure the link-local address, either automatically or manually.

Auto-MDI/MDIX Feature

For RJ-45 interfaces, the default auto-negotiation setting also includes the Auto-MDI/MDIX feature. Auto-MDI/MDIX eliminates the need for crossover cabling by performing an internal crossover when a straight cable is detected during the auto-negotiation phase. Either the speed or duplex must be set to auto-negotiate to enable Auto-MDI/MDIX for the interface. If you explicitly set both the speed and duplex to a fixed value, thus disabling auto-negotiation for both settings, then Auto-MDI/MDIX is also disabled. For Gigabit Ethernet, when the speed and duplex are set to 1000 and full, then the interface always auto-negotiates; therefore Auto-MDI/MDIX is always enabled and you cannot disable it.

Guidelines and Limitations for Interfaces

The following topics cover some of the limitations for interfaces.

Limitations for Interface Configuration

When you use the FDM to configure the device, there are several limitations to interface configuration. If you need any of the following features, you must use the FMC to configure the device.

- Routed firewall mode only is supported. You cannot configure transparent firewall mode interfaces.
- You cannot configure passive or ERSPAN interfaces.
- You cannot configure interfaces to be inline (in an inline set), or inline tap, for IPS-only processing. IPS-only mode interfaces bypass many firewall checks and only support IPS security policy. In comparison, Firewall mode interfaces subject traffic to firewall functions such as maintaining flows, tracking flow states at both IP and TCP layers, IP defragmentation, and TCP normalization. You can also optionally configure IPS functions for this firewall mode traffic according to your security policy.
- You cannot configure EtherChannel or redundant interfaces.
- You can only add one bridge group.
- You cannot configure PPPoE for IPv4. If the Internet interface is connected to a DSL, cable modem, or other connection to your ISP, and your ISP uses PPPoE to provide your IP address, you must use the FMC instead of the FDM.

- For the ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X, you can install an optional network interface module. Modules are only discovered during bootstrap (that is, during installation, when switching between local/remote management, and during a major/minor release upgrade, but not patch or hot fix upgrades). For a module that includes SFP interfaces, the FDM sets the speed and duplex to auto; however, the SFP interfaces do not support the speed and duplex set to auto. For these interfaces, select the right speed (for example, 1000), or select **Default** for the speed and duplex. The Default setting tells FDM to simply not configure the options, and thus leave them at their default settings (any existing configuration is not cleared). Please refer to the network module documentation to determine the maximum speed supported by the interface. You can also select **No Negotiate** for the speed setting if the interface accepts it, but select this option only if you are certain it is supported.



Note If you make a mistake and need to unconfigure **No Negotiate**, set the option to **Auto** and deploy. The deployment will fail. You can then set the option to **Default** and deploy again, and this should result in a successful deployment.

Maximum Number of VLAN Subinterfaces by Device Model

The device model limits the maximum number of VLAN subinterfaces that you can configure. Note that you can configure subinterfaces on data interfaces only, you cannot configure them on the management interface.

The following table explains the limits for each device model.


Model	Maximum VLAN Subinterfaces
ASA 5506-X	30
ASA 5506W-X	
ASA 5506H-X	
ASA 5508-X	50
ASA 5512-X	100
ASA 5515-X	
ASA 5516-X	100
ASA 5525-X	200
ASA 5545-X	300
ASA 5555-X	500

Configure a Physical Interface

At minimum, you must enable a physical interface to use it. You would also typically name it and configure IP addressing. You would not configure IP addressing if you intend to create VLAN subinterfaces, or if you intend to add the interface to a bridge group.

You can disable an interface to temporarily prevent transmission on the connected network. You do not need to remove the interface's configuration.

Procedure

- Step 1** Click **Device**, then click the link in the **Interfaces** summary.
- The interface list shows the available interfaces, their names, addresses, and states.
- Step 2** Click the edit icon () for the physical interface you want to edit.
- Step 3** Set the following:

Ethernet1/2
Edit Physical Interface

Interface Name: Mode: Status:

Most features work with named interfaces only, although some require unnamed interfaces.

Description:

IPv4 Address | IPv6 Address | Advanced

Type:


IP Address and Subnet Mask: /
e.g. 192.168.5.15/17 or 192.168.5.15/255.255.128.0

Standby IP Address and Subnet Mask: /
e.g. 192.168.5.16

- a) Set the **Interface Name**.

Set the name for the interface, up to 48 characters. Alphabetic characters must be lower case. For example, **inside** or **outside**. Without a name, the rest of the interface configuration is ignored. Unless you configure subinterfaces, the interface should have a name.

Note If you change the name, the change is automatically reflected everywhere you used the old name, including security zones, syslog server objects, and DHCP server definitions. However, you cannot remove the name until you first remove all configurations that use the name, because you typically cannot use an unnamed interface for any policy or setting.

- b) Set the **Status** slider to the enabled setting ()

If you intend to configure subinterfaces for this physical interface, you are probably done. Click **Save** and continue with [Configure VLAN Subinterfaces and 802.1Q Trunking, on page 110](#). Otherwise, continue.

Note Even when configuring subinterfaces, it is valid to name the interface and supply IP addresses. This is not the typical setup, but if you know that is what you need, you can configure it.

- c) (Optional) Set the **Description**.

The description can be up to 200 characters on a single line, without carriage returns.

Step 4 Click the **IPv4 Address** tab and configure the IPv4 address.

Select one of the following options from the **Type** field:

- **DHCP**—Choose this option if the address should be obtained from the DHCP server on the network. Change the following options if necessary:
 - **Route Metric**—If you obtain the default route from the DHCP server, the administrative distance to the learned route, between 1 and 255. The default is 1.
 - **Obtain Default Route**—Whether to get the default route from the DHCP server. You would normally select this option, which is the default.
- **Static**—Choose this option if you want to assign an address that should not change. Type in the interface's IP address and the subnet mask for the network attached to the interface. For example, if you attach the 10.100.10.0/24 network, you could enter 10.100.10.1/24. Ensure that the address is not already used on the network.

Note For an existing interface, your ability to change the address is constrained if you have a DHCP server configured for the interface. The new IP address must be on the same subnet as the DHCP address pool, and it cannot be part of that pool. If you need to configure an address on a different subnet, first delete the DHCP server configuration. See [Configuring the DHCP Server, on page 284](#).

Step 5 (Optional.) Click the **IPv6 Address** tab and configure the IPv6 address.

- **State**—To enable IPv6 processing and to automatically configure the link-local address when you do not configure the global address, select **Enabled**. The link local address is generated based on the interface MAC addresses (*Modified EUI-64* format).

Note Disabling IPv6 does not disable IPv6 processing on an interface that is configured with an explicit IPv6 address or that is enabled for autoconfiguration.

- **Address Auto Configuration**—Select this option to have the address automatically configured. IPv6 stateless autoconfiguration will generate a global IPv6 address only if the link on which the device resides has a router configured to provide IPv6 services, including the advertisement of an IPv6 global prefix for use on the link. If IPv6 routing services are not available on the link, you will get a link-local IPv6 address only, which you cannot access outside of the device's immediate network link. The link local address is based on the Modified EUI-64 interface ID.

Although RFC 4862 specifies that hosts configured for stateless autoconfiguration do not send Router Advertisement messages, the FTD device does send Router Advertisement messages in this case. Select **Suppress RA** to suppress messages and conform to the RFC.

- **Static Address/Prefix**—If you do not use stateless autoconfiguration, enter the full static global IPv6 address and network prefix. For example, 2001:0DB8::BA98:0:3210/48. For more information on IPv6 addressing, see [IPv6 Addressing, on page 101](#).

If you want to use the address as link local only, select the **Link - Local** option. Link local addresses are not accessible outside the local network. You cannot configure a link-local address on a bridge group interface.

Note A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. Note that we recommend automatically assigning the link-local address based on the Modified EUI-64 format. For example, if other devices enforce the use of the Modified EUI-64 format, then a manually-assigned link-local address may cause packets to be dropped.

- **Suppress RA**—Whether to suppress router advertisements. The FTD can participate in router advertisements so that neighboring devices can dynamically learn a default router address. By default, router advertisement messages (ICMPv6 Type 134) are periodically sent out each IPv6 configured interface.

Router advertisements are also sent in response to router solicitation messages (ICMPv6 Type 133). Router solicitation messages are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled router advertisement message.

You might want to suppress these messages on any interface for which you do not want the FTD device to supply the IPv6 prefix (for example, the outside interface).

Step 6 (Optional.) [Configure Advanced Options, on page 114](#).

The advanced settings have defaults that are appropriate for most networks. Edit them only if you are resolving network issues.

Step 7 Click **OK**.

What to do next

- Add the interfaces to the appropriate security zones. See [Configuring Security Zones, on page 89](#).

Configure Bridge Groups

A bridge group is a virtual interface that groups one or more interfaces. The main reason to group interfaces is to create a group of switched interfaces. Thus, you can attach workstations or other endpoint devices directly to the interfaces included in the bridge group. You do not need to connect them through a separate physical switch, although you can also attach a switch to a bridge group member.

The group members do not have IP addresses. Instead, all member interfaces share the IP address of the Bridge Virtual Interface (BVI). If you enable IPv6 on the BVI, member interfaces are automatically assigned unique link-local addresses.

You enable and disable the member interfaces individually. Thus, you can disable any unused interfaces without needing to remove them from the bridge group. The bridge group itself is always enabled.

You typically configure a DHCP server on the bridge group interface (BVI), which provides IP addresses for any endpoints connected through member interfaces. However, you can configure static addresses on the endpoints connected to the member interfaces if you prefer. All endpoints within the bridge group must have IP addresses on the same subnet as the bridge group IP address.

Guidelines and Limitations

- You can add one bridge group.
- For all ASA 5506-X models, on a new version 6.2+ system, or a reimaged 6.2+ system, the device comes pre-configured with bridge group BV11, named **inside**, which includes all data interfaces except for the **outside** interface. Thus, the device is pre-configured with one port used for linking to the Internet or other upstream network, and all other ports enabled and available for direct connections to endpoints. If you want to use an inside interface for a new subnet, you must first remove the needed interfaces from BV11.

Before you begin

Configure the interfaces that will be members of the bridge group. Specifically, each member interface must meet the following requirements:

- The interface must have a name.
- The interface cannot have any IPv4 or IPv6 addresses defined for it, either static or served through DHCP. If you need to remove the address from an interface that you are currently using, you might also need to remove other configurations for the interface, such as static routes, DHCP server, or NAT rules, that depend on the interface having an address.
- You must remove the interface from its security zone (if it is in a zone), and delete any NAT rules for the interface, before you can add it to a bridge group.

Procedure

Step 1 Click **Device**, then click the link in the **Interfaces** summary.

The interface list shows the available interfaces, their names, addresses, and states. If there is already a bridge group, it is a folder. Click the open/close arrow to view the member interfaces. Member interfaces also appear separately in the list.

Step 2 Do one of the following:

- Click the edit icon (🔗) for the BV11 bridge group.
- Select **Add Bridge Group Interface** from the gear drop-down list to create a new group.

Note You can have a single bridge group. If you already have a bridge group defined, you should edit that group instead of trying to create a new one. If you need to create a new bridge group, you must first delete the existing bridge group.
- Click the delete icon (🗑️) for the bridge group if you no longer need it. When you delete a bridge group, its members become standard routed interfaces, and any NAT rules or security zone membership are retained. You can edit the interfaces to give them IP addresses. If you want to add them to a new bridge group, first you need to remove the NAT rules and remove the interface from its security zone.

Step 3 Configure the following:

a) (Optional) Set the **Interface Name**.

Set the name for the bridge group, up to 48 characters. Alphabetic characters must be lower case. For example, **inside** or **outside**. Set the name if you want this BVI to participate in routing between it and other named interfaces.

Note If you change the name, the change is automatically reflected everywhere you used the old name, including security zones, syslog server objects, and DHCP server definitions. However, you cannot remove the name until you first remove all configurations that use the name, because you typically cannot use an unnamed interface for any policy or setting.

b) (Optional) Set the **Description**.

The description can be up to 200 characters on a single line, without carriage returns.

c) Edit the **Bridge Group Members** list.

You can add up to 64 interfaces or subinterfaces to a single bridge group.

- Add an interface—Click the plus icon (+), click one or more interfaces, and then click **OK**.
- Remove an interface—Mouse over an interface and click the **x** on the right side.

Step 4 Click the **IPv4 Address** tab and configure the IPv4 address.

Select one of the following options from the **Type** field:

- **Static**—Choose this option if you want to assign an address that should not change. Type in the bridge group's IP address and the subnet mask. All attached endpoints will be on this network. For models with a pre-configured bridge group, the default for the BV11 “inside” network is 192.168.1.1/24 (i.e. 255.255.255.0). Ensure that the address is not already used on the network.

Note For an existing bridge group, your ability to change the address is constrained if you have a DHCP server configured for the group. The new IP address must be on the same subnet as the DHCP address pool, and it cannot be part of that pool. If you need to configure an address on a different subnet, first delete the DHCP server configuration. See [Configuring the DHCP Server, on page 284](#).

- **Dynamic (DHCP)**—Choose this option if the address should be obtained from the DHCP server on the network. This is not the typical option for bridge groups, but you can configure it if needed. Change the following options if necessary:
 - **Route Metric**—If you obtain the default route from the DHCP server, the administrative distance to the learned route, between 1 and 255. The default is 1.
 - **Obtain Default Route**—Whether to get the default route from the DHCP server. You would normally select this option, which is the default.

Step 5 (Optional.) Click the **IPv6 Address** tab and configure the IPv6 address.

- **State**—To enable IPv6 processing and to automatically configure the link-local address when you do not configure the global address, select **Enabled**. The link local address is generated based on the interface MAC addresses (*Modified* EUI-64 format).

Note Disabling IPv6 does not disable IPv6 processing on an interface that is configured with an explicit IPv6 address or that is enabled for autoconfiguration.

- **Static Address/Prefix**—If you do not use stateless autoconfiguration, enter the full static global IPv6 address and network prefix. For example, 2001:0DB8::BA98:0:3210/48. For more information on IPv6 addressing, see [IPv6 Addressing, on page 101](#).

If you want to use the address as link local only, select the **Link - Local** option. Link local addresses are not accessible outside the local network. You cannot configure a link-local address on a bridge group interface.

Note A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. Note that we recommend automatically assigning the link-local address based on the Modified EUI-64 format. For example, if other devices enforce the use of the Modified EUI-64 format, then a manually-assigned link-local address may cause packets to be dropped.

- **Suppress RA**—Whether to suppress router advertisements. The FTD device can participate in router advertisements so that neighboring devices can dynamically learn a default router address. By default, router advertisement messages (ICMPv6 Type 134) are periodically sent out each IPv6 configured interface.

Router advertisements are also sent in response to router solicitation messages (ICMPv6 Type 133). Router solicitation messages are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled router advertisement message.

You might want to suppress these messages on any interface for which you do not want the FTD device to supply the IPv6 prefix (for example, the outside interface).

Step 6 (Optional.) [Configure Advanced Options, on page 114.](#)

You configure most advanced options on bridge group *member* interfaces, but some are available for the bridge group interface.

The advanced settings have defaults that are appropriate for most networks. Edit them only if you are resolving network issues.

Step 7 Click **OK**.

What to do next

- Ensure that all member interfaces that you intend to use are enabled.
- Configure a DHCP server for the bridge group. See [Configuring the DHCP Server, on page 284.](#)
- Add the member interfaces to the appropriate security zones. See [Configuring Security Zones, on page 89.](#)
- Ensure that policies, such as identity, NAT, and access, supply the required services for the bridge group and member interfaces.

Configure VLAN Subinterfaces and 802.1Q Trunking

VLAN subinterfaces let you divide a physical interface into multiple logical interfaces that are tagged with different VLAN IDs. An interface with one or more VLAN subinterfaces is automatically configured as an 802.1Q trunk. Because VLANs allow you to keep traffic separate on a given physical interface, you can increase the number of interfaces available to your network without adding additional physical interfaces or devices.

Create subinterfaces if you attach the physical interface to a trunk port on a switch. Create a subinterface for each VLAN that can appear on the switch trunk port. If you attach the physical interface to an access port on the switch, there is no point in creating a subinterface.

Guidelines and Limitations


- Preventing untagged packets on the physical interface—If you use subinterfaces, you typically do not also want the physical interface to pass traffic, because the physical interface passes untagged packets. Because the physical interface must be enabled for the subinterface to pass traffic, ensure that the physical interface does not pass traffic by not naming the interface. If you want to let the physical interface pass untagged packets, you can name the interface as usual.
- You cannot configure IP addresses on bridge group member interfaces, although you can modify advanced settings as needed.
- All subinterfaces on the same parent interface must be either bridge group members or routed interfaces; you cannot mix and match.
- FTD does not support the Dynamic Trunking Protocol (DTP), so you must configure the connected switch port to trunk unconditionally.
- You might want to assign unique MAC addresses to subinterfaces defined on the FTD device, because they use the same burned-in MAC address of the parent interface. For example, your service provider might perform access control based on the MAC address. Also, because IPv6 link-local addresses are generated based on the MAC address, assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the FTD device.


Procedure


Step 1 Click **Device**, then click the link in the **Interfaces** summary.

The interface list shows the available interfaces, their names, addresses, and states.

Step 2 Do one of the following:

- Select **Add Subinterface** from the gear drop-down list to create a new subinterface.
- Click the edit icon () for the subinterface you want to edit.

If you no longer need a subinterface, click the delete icon () for the subinterface to delete it.

Step 3 Set the **Status** slider to the enabled setting (.

Step 4 Configure the parent interface, name, and description:

a) Choose the **Parent Interface**.

The parent interface is the physical interface to which you want to add the subinterface. You cannot change the parent interface after you create the subinterface.

b) Set the **Subinterface Name**, up to 48 characters.

Alphabetic characters must be lower case. For example, **inside** or **outside**. Without a name, the rest of the interface configuration is ignored.

Note If you change the name, the change is automatically reflected everywhere you used the old name, including security zones, syslog server objects, and DHCP server definitions. However, you cannot remove the name until you first remove all configurations that use the name, because you typically cannot use an unnamed interface for any policy or setting.

c) (Optional) Set a **Description**.

The description can be up to 200 characters on a single line, without carriage returns.

d) Set the **VLAN ID**.

Enter the VLAN ID between 1 and 4094 that will be used to tag the packets on this subinterface.

e) Set the **Subinterface ID**.

Enter the subinterface ID as an integer between 1 and 4294967295. This ID is appended to the interface ID; for example Ethernet1/1.100. You can match the VLAN ID for convenience, but it is not required. You cannot change the ID after you create the subinterface.

Step 5 Click the **IPv4 Address** tab and configure the IPv4 address.

Select one of the following options from the **Type** field:

- **DHCP**—Choose this option if the address should be obtained from the DHCP server on the network. Change the following options if necessary:
 - **Route Metric**—If you obtain the default route from the DHCP server, the administrative distance to the learned route, between 1 and 255. The default is 1.
 - **Obtain Default Route**—Whether to get the default route from the DHCP server. You would normally select this option, which is the default.
- **Static**—Choose this option if you want to assign an address that should not change. Type in the interface's IP address and the subnet mask for the network attached to the interface. For example, if you attach the 10.100.10.0/24 network, you could enter 10.100.10.1/24. Ensure that the address is not already used on the network.

Note For an existing interface, your ability to change the address is constrained if you have a DHCP server configured for the interface. The new IP address must be on the same subnet as the DHCP address pool, and it cannot be part of that pool. If you need to configure an address on a different subnet, first delete the DHCP server configuration. See [Configuring the DHCP Server, on page 284](#).

Step 6 (Optional.) Click the **IPv6 Address** tab and configure the IPv6 address.

- **State**—To enable IPv6 processing and to automatically configure the link-local address when you do not configure the global address, select **Enabled**. The link local address is generated based on the interface MAC addresses (*Modified EUI-64* format).

Note Disabling IPv6 does not disable IPv6 processing on an interface that is configured with an explicit IPv6 address or that is enabled for autoconfiguration.

- **Address Auto Configuration**—Select this option to have the address automatically configured. IPv6 stateless autoconfiguration will generate a global IPv6 address only if the link on which the device resides has a router configured to provide IPv6 services, including the advertisement of an IPv6 global prefix for use on the link. If IPv6 routing services are not available on the link, you will get a link-local IPv6 address only, which you cannot access outside of the device's immediate network link. The link local address is based on the Modified EUI-64 interface ID.

Although RFC 4862 specifies that hosts configured for stateless autoconfiguration do not send Router Advertisement messages, the FTD device does send Router Advertisement messages in this case. Select **Suppress RA** to suppress messages and conform to the RFC.

- **Static Address/Prefix**—If you do not use stateless autoconfiguration, enter the full static global IPv6 address and network prefix. For example, 2001:0DB8::BA98:0:3210/48. For more information on IPv6 addressing, see [IPv6 Addressing, on page 101](#).

If you want to use the address as link local only, select the **Link - Local** option. Link local addresses are not accessible outside the local network. You cannot configure a link-local address on a bridge group interface.

Note A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. Note that we recommend automatically assigning the link-local address based on the Modified EUI-64 format. For example, if other devices enforce the use of the Modified EUI-64 format, then a manually-assigned link-local address may cause packets to be dropped.

- **Suppress RA**—Whether to suppress router advertisements. The FTD can participate in router advertisements so that neighboring devices can dynamically learn a default router address. By default, router advertisement messages (ICMPv6 Type 134) are periodically sent out each IPv6 configured interface.

Router advertisements are also sent in response to router solicitation messages (ICMPv6 Type 133). Router solicitation messages are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled router advertisement message.

You might want to suppress these messages on any interface for which you do not want the FTD device to supply the IPv6 prefix (for example, the outside interface).

Step 7 (Optional.) [Configure Advanced Options, on page 114](#).

The advanced settings have defaults that are appropriate for most networks. Edit them only if you are resolving network issues.

Step 8 Click **OK**.

What to do next

- Add the subinterfaces to the appropriate security zones. See [Configuring Security Zones, on page 89](#).

Configure Advanced Interface Options

Advanced options include setting the MTU, hardware settings, management only, MAC address, and other settings.

About the MTU

The MTU specifies the maximum frame *payload* size that the FTD device can transmit on a given Ethernet interface. The MTU value is the frame size *without* Ethernet headers, VLAN tagging, or other overhead. For example, when you set the MTU to 1500, the expected frame size is 1518 bytes including the headers, or 1522 when using VLAN. Do not set the MTU value higher to accommodate these headers.

Path MTU Discovery

The FTD device supports Path MTU Discovery (as defined in RFC 1191), which lets all devices in a network path between two hosts coordinate the MTU so they can standardize on the lowest MTU in the path.

MTU and Fragmentation

For IPv4, if an outgoing IP packet is larger than the specified MTU, it is fragmented into 2 or more frames. Fragments are reassembled at the destination (and sometimes at intermediate hops), and fragmentation can cause performance degradation. For IPv6, packets are typically not allowed to be fragmented at all. Therefore, your IP packets should fit within the MTU size to avoid fragmentation.

For UDP or ICMP, the application should take the MTU into account to avoid fragmentation.



Note The FTD device can receive frames larger than the configured MTU as long as there is room in memory.

MTU and Jumbo Frames

A larger MTU lets you send larger packets. Larger packets might be more efficient for your network. See the following guidelines:

- Matching MTUs on the traffic path—We recommend that you set the MTU on all FTD interfaces and other device interfaces along the traffic path to be the same. Matching MTUs prevents intermediate devices from fragmenting the packets.
- Accommodating jumbo frames—A jumbo frame is an Ethernet packet larger than the standard maximum of 1522 bytes (including Layer 2 header and VLAN header), up to 9216 bytes. You can set the MTU to 9000 bytes or higher to accommodate jumbo frames. The maximum depends on the model.



Note Increasing the MTU assigns more memory for jumbo frames, which might limit the maximum usage of other features, such as access rules. If you increase the MTU above the default 1500 on ASA 5500-X series devices, you must reboot the system.

Configure Advanced Options


Advanced interface options have default settings that are appropriate for most networks. Configure them only if you are resolving networking problems.

The following procedure assumes the interface is already defined. You can also edit these settings while initially editing or creating the interface.

Limitations

- For bridge groups, you configure most of these options on the member interfaces. Except for DAD attempts, these options are not available for the Bridge Virtual Interface (BVI).

Procedure

- Step 1** Click **Device**, then click the link in the **Interfaces** summary.
- The interface list shows the available interfaces, their names, addresses, and states.
- Step 2** Click the edit icon () for the interface you want to edit.
- Step 3** Click **Advanced Options**.
- Step 4** To make a data interface management only, select **Management Only**.
- A management only interface does not allow through traffic, so there is very little value in setting a data interface as management only. You cannot change this setting for the Management/Diagnostic interface, which is always management only.
- Step 5** Change the **MTU** (maximum transmission unit) to the desired value.
- The default MTU is 1500 bytes. The minimum and maximum depend on your platform. Set a high value if you typically see jumbo frames on your network.
- Note** If you increase MTU above 1500 on ASA 5500-X series devices, you must reboot the device.
- Step 6** (Physical interface only.) Modify the speed and duplex settings.
- The default is that the interface negotiates the best duplex and speed with the interface at the other end of the wire, but you can force a specific duplex or speed if necessary. Before setting these options for interfaces on a network module, please read [Limitations for Interface Configuration, on page 102](#).
- **Duplex**—Choose **Auto**, **Half** or **Full**. SFP interfaces only support **Full** duplex.
 - **Speed**—Choose a speed (varies depending on the model), or **Auto**.
- Step 7** Modify the **IPv6 Configuration** settings.
- **Enable DHCP for IPv6 address configuration**—Whether to set the Managed Address Configuration flag in the IPv6 router advertisement packet. This flag informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain addresses, in addition to the derived stateless autoconfiguration address.
 - **Enable DHCP for IPv6 non-address configuration**—Whether to set the Other Address Configuration flag in the IPv6 router advertisement packet. This flag informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain additional information from DHCPv6, such as the DNS server address.
 - **DAD Attempts**—How often the interface performs Duplicate Address Detection (DAD), from 0 - 600. The default is 1. During the stateless autoconfiguration process, DAD verifies the uniqueness of new unicast IPv6 addresses before the addresses are assigned to interfaces. If the duplicate address is the link-local address of the interface, the processing of IPv6 packets is disabled on the interface. If the duplicate address is a global address, the address is not used. The interface uses neighbor solicitation messages to perform Duplicate Address Detection. Set the value to 0 to disable duplicate address detection (DAD) processing.
- Step 8** Click **OK**.
-

Monitoring Interfaces

You can view some basic information about interfaces in the following areas:

- **Device.** Use the port graphic to monitor the current state of the interfaces. Mouse over a port to see its IP addresses and enabled and link statuses. The IP addresses can be statically assigned or obtained using DHCP.

Interface ports use the following color coding:

- Green—The interface is configured, enabled, and the link is up.
 - Gray—The interface is not enabled.
 - Orange/Red—The interface is configured and enabled, but the link is down. If the interface is wired, this is an error condition that needs correction. If the interface is not wired, this is the expected status.
- **Monitoring > System.** The **Throughput** dashboard shows information on traffic flowing through the system. You can view information on all interfaces, or you can select a specific interface to examine.
 - **Monitoring > Ingress Zones and Egress Zones.** These dashboards show statistics based on zones, which are composed of interfaces. You can drill into this information for more detail.

Monitoring Interfaces in the CLI

You can also log into the device CLI and use the following commands to get more detailed information about interface-related behavior and statistics.

- **show interface** displays interface statistics and configuration information. This command has many keywords you can use to get to the information you need. Use ? as a keyword to see the available options.
- **show ipv6 interface** displays IPv6 configuration information about the interfaces.
- **show bridge-group** displays information about Bridge Virtual Interfaces (BVI), including member information and IP addresses.
- **show conn** displays information about the connections currently established through the interfaces.
- **show traffic** displays statistics about traffic flowing through each interface.
- **show ipv6 traffic** displays statistics about IPv6 traffic flowing through the device.
- **show dhcpd** displays statistics and other information about DHCP usage on the interfaces, particularly about the DHCP servers configured on interfaces.

Examples for Interfaces

The use case chapter includes the following interface-related examples:

- [How to Configure the Device in FDM, on page 31](#)
- [How to Add a Subnet, on page 57](#)



CHAPTER 7

Routing

The system uses a routing table to determine the egress interface for packets entering the system. The following topics explain routing basics and how to configure routing on the device.

- [Routing Overview, on page 117](#)
- [Static Routes, on page 119](#)
- [Monitoring Routing, on page 121](#)

Routing Overview

The following topics describe how routing behaves within the FTD device. Routing is the act of moving information across a network from a source to a destination. Along the way, at least one intermediate node is typically encountered. Routing involves two basic activities: determining optimal routing paths and transporting packets through a network.

The Routing Table and Route Selection

When NAT translations (xlates) and rules do not determine the egress interface, the system uses the routing table to determine the path for a packet.

Routes in the routing table include a metric called “administrative distance” that provides a relative priority to a given route. If a packet matches more than one route entry, the one with the lowest distance is used. Directly connected networks (those defined on an interface) have the distance 0, so they are always preferred. Static routes have a default distance of 1, but you can create them with any distance between 1-254.

Routes that identify a specific destination take precedence over the default route (the route whose destination is 0.0.0.0/0 or ::/0).

How Forwarding Decisions Are Made

Forwarding decisions are made as follows:

- If the destination does not match an entry in the routing table, the packet is forwarded through the interface specified for the default route. If a default route has not been configured, the packet is discarded.
- If the destination matches a single entry in the routing table, the packet is forwarded through the interface associated with that route.

- If the destination matches more than one entry in the routing table, then the packet is forwarded out of the interface associated with the route that has the longer network prefix length.

For example, a packet destined for 192.168.32.1 arrives on an interface with the following routes in the routing table:

- 192.168.32.0/24 gateway 10.1.1.2
- 192.168.32.0/19 gateway 10.1.1.3

In this case, a packet destined to 192.168.32.1 is directed toward 10.1.1.2, because 192.168.32.1 falls within the 192.168.32.0/24 network. It also falls within the other route in the routing table, but 192.168.32.0/24 has the longest prefix within the routing table (24 bits versus 19 bits). Longer prefixes are always preferred over shorter ones when forwarding a packet.



Note Existing connections continue to use their established interfaces even if a new similar connection would result in different behavior due to a change in routes.

Routing Table for Management Traffic

As a standard security practice, it is often necessary to segregate and isolate management (from-the-device) traffic from data traffic. To achieve this isolation, the FTD device uses a separate routing table for management-only traffic vs. data traffic. Separate routing tables means that you can create separate default routes for data and management as well.

Types of Traffic for Each Routing Table

Through-the-device traffic always uses the data routing table.

From-the-device traffic, depending on the type, uses either the management-only routing table or the data routing table by default. If a match is not found in the default routing table, it checks the other routing table.

- Management-only table from-the-device traffic includes AAA server communications.
- Data table from-the-device traffic includes DNS server lookups and DDNS. An exception is if you only specify the Diagnostic interface for DNS, then the FTD device will only use the management-only table.

Interfaces Included in the Management-Only Routing Table

Management-only interfaces include any the Management x/x interfaces as well as any interfaces that you have configured to be management-only.



Note The Management virtual interface uses its own Linux routing table that is not part of the FTD route lookup. Traffic originating on the Management interface includes the FDM management sessions, licensing communication, and database updates. The Diagnostic logical interface, on the other hand, uses the management-only routing table described in this section.

Fallback to the Other Routing Table

If a match is not found in the default routing table, it checks the other routing table.

Using the Non-Default Routing Table

If you need from-the-box traffic to go out an interface that isn't in its default routing table, then you might need to specify that interface when you configure it, rather than relying on the fall back to the other table. The FTD will only check routes for the specified interface. For example, if you need to communicate with a RADIUS server on a data interface, then specify that interface in the RADIUS configuration. Otherwise, if there is a default route in the management-only routing table, then it will match the default route and never fall back to the data routing table.

Equal-Cost Multi-Path (ECMP) Routing

The FTD device supports Equal-Cost Multi-Path (ECMP) routing.

You can have up to 8 equal cost static or dynamic routes per interface. For example, you can configure multiple default routes on the outside interface that specify different gateways.

```
route for 0.0.0.0 0.0.0.0 through outside to 10.1.1.2
route for 0.0.0.0 0.0.0.0 through outside to 10.1.1.3
route for 0.0.0.0 0.0.0.0 through outside to 10.1.1.4
```

In this case, traffic is load-balanced on the outside interface between 10.1.1.2, 10.1.1.3, and 10.1.1.4. Traffic is distributed among the specified gateways based on an algorithm that hashes the source and destination IP addresses, incoming interface, protocol, source and destination ports.

Static Routes

You can create static routes to provide basic routing for your network.

About Static and Default Routes

To route traffic to a non-connected host or network, you must define a route to the host or network, either using static or dynamic routing. Generally, you must configure at least one static route: a default route for all traffic that is not routed by other means to a default network gateway, typically the next hop router.

Default Route

The simplest option is to configure a default static route to send all traffic to an upstream router, relying on the router to route the traffic for you. A default route identifies the gateway IP address to which the FTD device sends all IP packets for which it does not have a learned or static route. A default static route is simply a static route with 0.0.0.0/0 (IPv4) or ::/0 (IPv6) as the destination IP address.

You should always define a default route.

Because the FTD device uses separate routing tables for data traffic and for management traffic, you can optionally configure a default route for data traffic and another default route for management traffic. Note that from-the-device traffic uses either the management-only or data routing table by default depending on the type, but will fall back to the other routing table if a route is not found. Default routes will always match traffic, and will prevent a fall back to the other routing table. In this case, you must specify the interface you want to use for egress traffic if that interface is not in the default routing table. The Diagnostic interface is included in the management-only table. The special Management interface uses a separate Linux routing table, and has its own default route.

Static Routes

You might want to use static routes in the following cases:

- Your networks use an unsupported router discovery protocol.
- Your network is small and you can easily manage static routes.
- You do not want the traffic or CPU overhead associated with routing protocols.
- In some cases, a default route is not enough. The default gateway might not be able to reach the destination network, so you must also configure more specific static routes. For example, if the default gateway is outside, then the default route cannot direct traffic to any inside networks that are not directly connected to the FTD device.
- You are using a feature that does not support dynamic routing protocols.

Guidelines for Static Routing

Bridge Groups

- In routed mode, you must specify the BVI as the gateway; you cannot specify the member interface.
- For traffic that originates on the FTD device (such as syslog or SNMP) that is destined through a bridge group member interface for a non-directly connected network, you need to configure either a default route or static routes so the FTD device knows out of which bridge group member interface to send traffic. If you have servers that cannot all be reached through a single default route, then you must configure static routes.

Configuring Static Routes

Define static routes to tell the system where to send packets that are not bound for networks that are directly connected to the interfaces on the system.

You need at least one static route, the default route, for network 0.0.0.0/0. This route defines where to send packets whose egress interface cannot be determined by existing NAT xlates (translations) or static NAT rules, or other static routes.

You might need other static routes if the default gateway cannot be used to get to all networks. For example, the default route is usually an upstream router on the outside interface. If there are additional inside networks that are not directly connected to the device, and they cannot be accessed through the default gateway, you need static routes for each of those inside networks.

You cannot define static routes for the networks that are directly connected to system interfaces. The system automatically creates these routes.

Procedure

Step 1 Click **Device**, then click the link in the **Routing** summary.

Step 2 On the **Static Routing** page, do one of the following:

- To add a new route, click +.

- Click the edit icon (✎) for the route you want to edit.

If you no longer need a route, click the trash can icon for the route to delete it.

Step 3 Configure the route properties

- **Name**—A display name for the route.
- **Description**—An optional description of the purpose for the route.
- **Interface**—Select the interface through which you want to send traffic. The gateway address needs to be accessible through this interface.

For bridge groups, you configure the route for the bridge group interface (BVI), not for the member interfaces.

- **Protocol**—Select whether the route is for an **IPv4** or **IPv6** address.
- **Networks**—Select the network objects that identify the destination networks or hosts that should use the gateway in this route.

To define a default route, use the pre-defined any-ipv4 or any-ipv6 network objects, or create an object for the 0.0.0.0/0 (IPv4) or ::/0 (IPv6) network.

- **Gateway**—Select the host network object that identifies the IP address for the gateway. Traffic is sent to this address. You cannot use the same gateway for routes on more than one interface.
- **Metric**—The administrative distance for the route, between 1 and 254. The default for static routes is 1. If there are additional routers between the interface and the gateway, enter the number of hops as the administrative distance.

Administrative distance is a parameter used to compare routes. The lower the number, the higher precedence the route is given. Connected routes (networks directly connected to an interface on the device) always take precedence over static routes.

Step 4 Click OK.

Monitoring Routing

To monitor and troubleshoot routing, log into the device CLI and use the following commands.

- **show route** displays the routing table for the data interfaces, including routes for directly-connected networks.
- **show ipv6 route** displays the IPv6 routing table for the data interfaces, including routes for directly-connected networks.
- **show network** displays the configuration for the virtual Management interface, including the management gateway. Routing through the virtual Management interface is not handled by the data interface routing table, unless you specify data-interfaces as the management gateway.
- **show network-static-routes** displays static routes configured for the virtual Management interface using the **configure network static-routes** command. Normally, there will not be any static routes, as the

management gateway suffices for management routing in most cases. These routes are not available to traffic on the data interfaces. This command is not available in the CLI console.



PART II

Security Policies

- [Identity Policies, on page 125](#)
- [Access Control, on page 139](#)
- [Network Address Translation \(NAT\), on page 161](#)



CHAPTER 8

Identity Policies

You can use identity policies to collect user identity information from connections. You can then view usage based on user identity in the dashboards, and configure access control based on user or user group.

- [Identity Policy Overview, on page 125](#)
- [Configuring Identity Policies, on page 128](#)
- [Enabling Transparent User Authentication, on page 134](#)
- [Monitoring Identity Policies, on page 137](#)
- [Examples for Identity Policies, on page 137](#)

Identity Policy Overview

You can use identity policies to detect the user who is associated with a connection. By identifying the user, you can correlate threat, endpoint, and network intelligence with user identity information. By linking network behavior, traffic, and events directly to individual users, the system can help you identify the source of policy breaches, attacks, or network vulnerabilities.

For example, you can identify who owns the host targeted by an intrusion event, and who initiated an internal attack or port scan. You can also identify high bandwidth users and users who are accessing undesirable web sites or applications.

User detection goes beyond collecting data for analysis. You can also write access rules based on user name or user group name, selectively allowing or blocking access to resources based on user identity.

Establishing User Identity through Active Authentication

Authentication is the act of confirming the identity of a user.

With active authentication, when an HTTP traffic flow comes from an IP address for which the system has no user-identity mapping, you can decide whether to authenticate the user who initiated the traffic flow against the directory configured for the system. If the user successfully authenticates, the IP address is considered to have the identity of the authenticated user.

Failure to authenticate does not prevent network access for the user. Your access rules ultimately decide what access to provide these users.

Limitations on Number of Users

FDM can download information on up to 2000 users from the directory server.

If your directory server includes more than 2000 user accounts, you will not see all possible names when selecting users in an access rule or when viewing user-based dashboard information. You can write rules on only those names that were downloaded.

The limit also applies to the names associated with groups. If a group has more than 2000 members, only the 2000 names that were downloaded can be matched against the group membership.

If you have more than 2000 users, consider using the FMC (the remote manager) instead of FDM. FMC supports significantly more users.

Supported Directory Servers

You can use Microsoft Active Directory (AD) on Windows Server 2008 and 2012.

Note the following about your server configuration:

- If you want to perform user control on user groups or on users within groups, you must configure user groups on the directory server. The system cannot perform user group control if the server organizes the users in basic object hierarchy.
- The directory server must use the field names listed in the following table in order for the system to retrieve user metadata from the servers for that field.

Metadata	Active Directory Field
LDAP user name	samaccountname
first name	givenname
last name	sn
email address	mail userprincipalname (if mail has no value)
department	department distinguishedname (if department has no value)
telephone number	telephonenumber

Determining the Directory Base DN

When you configure directory properties, you need to specify the common base distinguished name (DN) for users and groups. The base is defined in your directory server, and differs from network to network. You must enter the correct bases for identity policies to work. If the base is wrong, the system cannot determine user or group names, and thus identity-based policies will be inoperable.



Tip To get the correct bases, consult the administrator who is responsible for the directory servers.

For active directory, you can determine the correct bases by logging into the Active Directory server as domain administrator, and using the **dsquery** command at a command prompt as follows to determine the bases:

User search base

Enter the **dsquery user** command with a known username (partial or complete) to determine the base distinguished name. For example, the following command uses the partial name “John*” to return information for all users that start with “John.”

```
C:\Users\Administrator>dsquery user -name "John*"
"CN=John Doe,CN=Users,DC=csc-lab,DC=example,DC=com"
```

The base DN would be “DC=csc-lab,DC=example,DC=com.”

Group search base

Enter the **dsquery group** command with a known group name to determine the base distinguished name. For example, the following command uses the group name Employees to return the distinguished name:

```
C:\>dsquery group -name "Employees"
"CN=Employees,CN=Users,DC=csc-lab,DC=example,DC=com"
```

The group base DN would be “DC=csc-lab,DC=example,DC=com.”

You can also use the ADSI Edit program to browse the Active Directory structure (**Start > Run > adsiedit.msc**). In ADSI Edit, right click any object, such as an organizational unit (OU), group, or user, and choose **Properties** to view the distinguished name. You can then copy the string of DC values as the base.

To verify that you have the correct base:

1. Click the Test Connection button in the directory properties to verify connectivity. Resolve any problems, and save the directory properties.
2. Commit changes to the device.
3. Create an access rule, select the **Users** tab, and try to add known user and group names from the directory. You should see auto-complete suggestions as you type for matching users and groups in the realm that contains the directory. If these suggestions appear in a drop-down list, then the system was able to query the directory successfully. If you see no suggestions, and you are certain the string you typed should appear in a user or group name, you need to correct the corresponding search base.

Dealing with Unknown Users

When you configure the directory server for the identity policy, the system downloads user and group membership information from the directory server. This information is refreshed every 24 hours at midnight, or whenever you edit and save the directory configuration (even if you do not make any changes).

If a user succeeds in authenticating when prompted by an active authentication identity rule, but the user’s name is not in the downloaded user identity information, the user is marked as Unknown. You will not see the user’s ID in identity-related dashboards, nor will the user match group rules.

However, any access control rules for the Unknown user will apply. For example, if you block connections for Unknown users, these users are blocked even though they succeeded in authenticating (meaning that the directory server recognizes the user and the password is valid).

Thus, when you make changes to the directory server, such as adding or deleting users, or changing group membership, these changes are not reflected in policy enforcement until the system downloads the updates from the directory.

If you do not want to wait until the daily midnight update, you can force an update by editing the directory server information (from **Policies > Identity**, click the **Directory Server** button). Click **Save**, then deploy changes. The system will immediately download the updates.



Note You can check whether new or deleted user information is on the system by going to **Policies > Access Control**, clicking the **Add Rule (+)** button, and looking at the list of users on the **Users** tab. If you cannot find a new user, or you can find a deleted user, then the system has old information.

Configuring Identity Policies

You can use identity policies to collect user identity information from connections. You can then view usage based on user identity in the dashboards, and configure access control based on user or user group.

The following is an overview of how to configure the elements required to obtain user identity through identity policies.

Procedure



Step 1 Select **Policies > Identity**.



If you have not yet defined an identity policy, you are prompted to start a wizard to configure it. Click **Get Started** to start the wizard. The wizard walks you through the following steps:

- a) [Configure Directory Server, on page 129](#)
- b) [Configure the Active Authentication Captive Portal, on page 130](#)

Step 2 Manage the identity policy.

After you configure identity settings, this page lists all rules in order. Rules are matched against traffic from top to bottom with the first match determining the action to apply. You can do the following from this page:

- To enable or disable the identity policy, click the **Identity Policy** toggle.
- To change the directory server configuration, click the **Directory Server** button (.
- To change the active authentication captive portal configuration, click the **Active Authentication** button (.
- To move a rule, edit it and select the new location from the **Order** drop-down list.
- To configure rules:
 - To create a new rule, click the + button.

- To edit an existing rule, click the edit icon () for the rule (in the Actions column). You can also selectively edit a rule property by clicking on the property in the table.
- To delete a rule you no longer need, click the delete icon () for the rule (in the Actions column).

For more information on creating and editing identity rules, see [Configure Identity Rules, on page 131](#).

Configure Directory Server

The directory server contains information about the users and user groups who are allowed access to your network. The system downloads updated information about all users and groups every day in the last hour of the day (UTC).

Work with your directory administrator to get the values required to configure the directory server properties.




Note After you add the realm, you can verify your settings and test the connection by clicking the **Directory Server** button and then clicking the **Test** button in the Directory Server dialog box. If the test fails, verify all fields and ensure there is a network path between the management IP address and the directory server.

Procedure

Step 1 Select **Policies > Identity**.

Step 2 Do one of the following:

- If you have not configured the directory or identity rules yet, click **Get Started** to start the Identity Policy wizard. You are first prompted to configure the directory server.
- Click the **Directory Server** button (.

Step 3 Fill in the following information about your directory server:

- **Name**—A name for the directory realm.
- **Type**—The type of directory server. Active Directory is the only supported type, and you cannot change this field.
- **Directory Username, Directory Password**—The distinguished username and password for a user with appropriate rights to the user information you want to retrieve. For Active Directory, the user does not need elevated privileges. You can specify any user in the domain. The username must be fully qualified; for example, Administrator@example.com (not simply Administrator).
- **Base DN**—The directory tree for searching or querying user and group information, that is, the common parent for users and groups. For example, dc=example,dc=com. For information on finding the base DN, see [Determining the Directory Base DN, on page 126](#).
- **AD Primary Domain**— The fully qualified Active Directory domain name that the device should join. For example, example.com.

- **Hostname/IP Address**—The hostname or IP address of the directory server. If you use an encrypted connection to the server, you must enter the fully-qualified domain name, not the IP address.
- **Port**—The port number used for communications with the server. The default is 389. Use port 636 if you select LDAPS as the encryption method.
- **Encryption**—To use an encrypted connection for downloading user and group information, select the desired method, **STARTTLS** or **LDAPS**. The default is **None**, which means that user and group information is downloaded in clear text.
 - **STARTTLS** negotiates the encryption method, and uses the strongest method supported by the directory server. Use port 389.
 - **LDAPS** requires LDAP over SSL. Use port 636.
- **SSL Certificate**—If you select an encryption method, upload a CA certificate to enable a trusted connection between the system and the directory server. If you are using a certificate to authenticate, the name of the server in the certificate must match the server Hostname / IP Address. For example, if you use 10.10.10.250 as the IP address but ad.example.com in the certificate, the connection fails.

Step 4 Click **Next** (in the wizard) or **Save**.

Configure the Active Authentication Captive Portal

When an identity rule requires active authentication for a user, the user is redirected to the captive portal port on the interface through which they are connected and then they are prompted to authenticate. If you do not upload a certificate, users are presented with a self-signed certificate. Users will have to accept the certificate if you do not upload a certificate that their browsers already trust.



Note For the HTTP Basic, HTTP Response Page, and NTLM authentication methods, the user is redirected to the captive portal using the IP address of the interface. However, for HTTP Negotiate, the user is redirected using the fully-qualified DNS name *firewall-hostname.AD-domain-name*. If you want to use HTTP Negotiate, you must also update your DNS server to map this name to the IP addresses of all inside interfaces where you are requiring active authentication. Otherwise, the redirection cannot complete, and users cannot authenticate.

Before you begin

Ensure that time settings are consistent among the directory servers, FTD device, and clients. A time shift among these devices can prevent successful user authentication. "Consistent" means that you can use different time zones, but the time should be the same relative to those zones; for example, 10 AM PST = 1 PM EST.

Procedure

Step 1 Select **Policies > Identity**.

Step 2 Do one of the following:

- If you are using the **Get Started** wizard, click **Next** after configuring the directory server.
- Click the **Active Authentication** button ()

Step 3 Configure the following options:

- **Server Certificate**—The CA certificate to present to users during active authentication. The certificate must be an X509 certificate in PEM or DER format. Paste in the certificate, or click **Upload Certificate** and select the certificate file. The default is to present a self-signed certificate during user authentication.
- **Certificate Key**—The key for the server certificate. Paste in the key, or click **Upload Key** and select the key file.
- **Port**—The captive portal port. The default is 885 (TCP). If you configure a different port, it must be in the range 1025-65535.

Step 4 Click **Save**.

Configure Identity Rules

Identity rules determine whether user identity information should be collected for matching traffic. You can configure No Authentication if you do not want to get user identity information for matching traffic.

Keep in mind that regardless of your rule configuration, active authentication is performed on HTTP traffic only. Thus, you do not need to create rules to exclude non-HTTP traffic from active authentication. You can simply apply an active authentication rule to all sources and destinations if you want to get user identity information for all HTTP traffic.





Note Also keep in mind that a failure to authentication has no impact on network access. Identity policies collect user identity information only. You must use access rules if you want to prevent users who failed to authenticate from accessing the network.

Procedure

Step 1 Select **Policies > Identity**.

Step 2 Do any of the following:

- To create a new rule, click the + button.
- To edit an existing rule, click the edit icon () for the rule.

To delete a rule you no longer need, click the delete icon () for the rule.

Step 3 In **Order**, select where you want to insert the rule in the ordered list of rules.

Rules are applied on a first-match basis, so you must ensure that rules with highly specific traffic matching criteria appear above policies that have more general criteria that would otherwise apply to the matching traffic.

The default is to add the rule to the end of the list. If you want to change a rule's location later, edit this option.

Step 4 Select the type of **User Authentication**.

You must select the AD identity realm that includes the user accounts for passive and active authentication rules.

- **Active**—Use active authentication to determine user identity. Active authentication is applied to HTTP traffic only. If any other type of traffic matches an identity policy that requires or allows active authentication, then active authentication will not be attempted.
- **No Auth**—Do not obtain user identity. Identity-based access rules will not be applied to this traffic. These users are marked as **No Authentication Required**.

Step 5 (Active Authentication only.) Select the authentication method (**Type**) supported by your directory server.

- **HTTP Basic**—Authenticate users using an unencrypted HTTP Basic Authentication (BA) connection. Users log in to the network using their browser's default authentication popup window. This is the default.
- **NTLM**—Authenticate users using an NT LAN Manager (NTLM) connection. This selection is only available when you select an AD realm. Users log in to the network using their browser's default authentication popup window, although you can configure IE and Firefox browsers to transparently authenticate using their Windows domain login (see [Enabling Transparent User Authentication, on page 134](#)).
- **HTTP Negotiate**—Allow the device to negotiate the method between the user agent (the application the user is using to initiate the traffic flow) and the Active Directory server. Negotiation results in the strongest commonly supported method being used, in order, NTLM, then basic. Users log in to the network using their browser's default authentication popup window.
- **HTTP Response Page**—Prompt users to authenticate using a system-provided web page. This is a form of HTTP Basic authentication.

Note For the HTTP Basic, HTTP Response Page, and NTLM authentication methods, the user is redirected to the captive portal using the IP address of the interface. However, for HTTP Negotiate, the user is redirected using the fully-qualified DNS name *firewall-hostname.AD-domain-name*. If you want to use HTTP Negotiate, you must also update your DNS server to map this name to the IP addresses of all inside interfaces where you are requiring active authentication. Otherwise, the redirection cannot complete, and users cannot authenticate.

Step 6 (Active authentication only.) Select **Fall Back as Guest > On/Off** to determine whether users who fail active authentication are labeled as Guest users.

Users get 3 chances to successfully authenticate. If they fail, your selection for this option determines how the user is marked. You can write access rules based on these values.

- **Fall Back as Guest > On**—Users are marked as **Guest**.
- **Fall Back as Guest > Off**—Users are marked as **Failed Authentication**.

Step 7 Define the traffic matching criteria on the **Source/Destination** tab.

Keep in mind that active authentication will be attempted with HTTP traffic only. Therefore, there is no need to configure No Auth rules for non-HTTP traffic, and there is no point in creating Active Authentication rules for any non-HTTP traffic.

The Source/Destination criteria of an identity rule define the security zones (interfaces) through which the traffic passes, the IP addresses or the country or continent (geographical location) for the IP address, or the protocols and ports used in the traffic. The default is any zone, address, geographical location, protocol, and port.

To modify a condition, you click the + button within that condition, select the desired object or element, and click **OK** in the popup dialog box. If the criterion requires an object, you can click **Create New Object** if the object you require does not exist. Click the **x** for an object or element to remove it from the policy.

You can configure the following traffic matching criteria.

Source Zones, Destination Zones

The security zone objects that define the interfaces through which the traffic passes. You can define one, both, or neither criteria: any criteria not specified applies to traffic on any interface.

- To match traffic leaving the device from an interface in the zone, add that zone to the **Destination Zones**.
- To match traffic entering the device from an interface in the zone, add that zone to the **Source Zones**.
- If you add both source and destination zone conditions to a rule, matching traffic must originate from one of the specified source zones and egress through one of the destination zones.

Use this criteria when the rule should apply based on where the traffic enters or exits the device. For example, if you want to ensure that user identity is collected from all traffic originating from inside networks, select an inside zone as the **Source Zones** while leaving the destination zone empty.

Source Networks, Destination Networks

The network objects or geographical locations that define the network addresses or locations of the traffic.

- To match traffic from an IP address or geographical location, configure the **Source Networks**.
- To match traffic to an IP address or geographical location, configure the **Destination Networks**.
- If you add both source and destination network conditions to a rule, matching traffic must originate from one of the specified IP addresses and be destined for one of the destination IP addresses.

When you add this criteria, you select from the following tabs:

- **Network**—Select the network objects or groups that define the source or destination IP addresses for the traffic you want to control.
- **Geolocation**—Select the geographical location to control traffic based on its source or destination country or continent. Selecting a continent selects all countries within the continent. Besides selecting geographical location directly in the rule, you can also select a geolocation object that you created to define the location. Using geographical location, you could easily restrict access to a particular country without needing to know all of the potential IP addresses used there.

Note To ensure you are using up-to-date geographical location data to filter your traffic, Cisco strongly recommends that you regularly update the geolocation database (GeoDB).

Source Ports, Destination Ports/Protocols

The port objects that define the protocols used in the traffic. For TCP/UDP, this can include ports.

- To match traffic from a protocol or port, configure the **Source Ports**. Source ports can be TCP/UDP only.
- To match traffic to a protocol or port, configure the **Destination Ports/Protocols**.
- To match traffic both originating from specific TCP/UDP ports and destined for specific TCP/UDP ports, configure both. If you add both source and destination ports to a condition, you can only add ports that

share a single transport protocol, TCP or UDP. For example, you could target traffic from port TCP/80 to port TCP/8080.

Step 8 Click OK.

Enabling Transparent User Authentication

If you configure the identity policy to allow for active authentication, you can use the following authentication methods to acquire user identity:

HTTP Basic

With HTTP basic authentication, users are always prompted to authenticate with their directory username and password. The password is transmitted in clear text. For that reason, basic authentication is not considered a secure form of authentication.

Basic is the default authentication mechanism.

HTTP Response Page

This is a type of HTTP basic authentication, where the user is presented with a login browser page.

NTLM, HTTP Negotiate (Integrated Windows Authentication for Active Directory)

With integrated Windows authentication, you take advantage of the fact that users log into a domain to use their workstation. The browser tries to use this domain login when accessing a server, including the FTD captive portal during active authentication. The password is not transmitted. If authentication is successful, the user is transparently authenticated; the user is unaware that any authentication challenge was made or satisfied.

If the browser cannot satisfy an authentication request using the domain login credentials, the user is prompted for username and password, which is the same user experience as basic authentication. Thus, if you configure integrated Windows authentication, it can reduce the need for users to supply credentials when accessing the network or servers in the same domain.

Note that HTTP Negotiate picks the strongest method supported by both the Active directory server and the user agent. If negotiation selects HTTP Basic as the authentication method, you will not get transparent authentication. The order of strength is NTLM, then basic. Negotiation must select NTLM for transparent authentication to be possible.

You must configure client browsers to support integrated Windows authentication to enable transparent authentication. The following sections explain the general requirements and basic configuration of integrated Windows authentication for some commonly used browsers that support it. Users should consult the help for their browser (or other user agent) for more detailed information, because the techniques can change between software releases.



Tip Not all browsers support integrated Windows authentication, such as Chrome and Safari (based on the versions available when this was written). Users will be prompted for username and password. Consult the browser's documentation to determine if support is available in the version you use.

Requirements for Transparent Authentication

Users must configure their browser or user agent to implement transparent authentication. They can do this individually, or you can configure it for them and push the configuration to client workstations using your software distribution tools. If you decide to have users do it themselves, ensure that you provide the specific configuration parameters that work for your network.

Regardless of browser or user agent, you must implement the following general configuration:

- Add the FTD interface through which users connect to the network to the Trusted Sites list. You can use the IP address or if available, the fully-qualified domain name (for example, `inside.example.com`). You can also use wildcards or partial addresses to create a generalized trusted site. For example, you can typically cover all internal sites using `*.example.com` or simply `example.com`, trusting all servers in your network (use your own domain name). If you add the specific address of the interface, you might need to add several addresses to the trusted sites to account for all user access points to the network.
- Integrated Windows authentication does not work through a proxy server. Therefore, you must either not use a proxy, or you must add the FTD interface to the addresses excluded from going through the proxy. If you decide that you must use a proxy, users will be prompted for authentication even if you use NTLM.



Tip Configuring transparent authentication is not a requirement, but a convenience to end users. If you do not configure transparent authentication, users are presented with a login challenge for all authentication methods.

Configuring Internet Explorer for Transparent Authentication

To configure Internet Explorer for NTLM transparent authentication:

Procedure

- Step 1** Select **Tools > Internet Options**.
- Step 2** Select the **Security** tab, select the **Local Intranet** zone, then do the following:
- a) Click the **Sites** button to open the list of trusted sites.
 - b) Ensure that at least one of the following options is selected:
 - **Automatically detect intranet network**. If you select this option, all other options are disabled.
 - **Include all sites that bypass the proxy**.
 - c) Click **Advanced** to open the Local Intranet Sites dialog box, then paste the URL you want to trust into the **Add Site** box and click **Add**.
Repeat the process if you have more than one URL. Use wildcards to specify a partial URL, such as `http://*.example.com` or simply `*.example.com`.
Close the dialog boxes to return to the Internet Options dialog box.
 - d) With **Local Intranet** still selected, click **Custom Level** to open the Security Settings dialog box. Find the **User Authentication > Logon** setting and select **Automatic logon only in Intranet zone**. Click **OK**.

Step 3 In the Internet Options dialog box, click the **Connections** tab, then click **LAN Settings**.

If **Use a proxy server for your LAN** is selected, you need to ensure that the FTD interface bypasses the proxy. Do any of the following as appropriate:

- Select **Bypass proxy server for local addresses**.
- Click **Advanced** and enter the address into the **Do not use proxy server for addresses beginning with** box. You can use wildcards, for example, ***.example.com**.

Configuring Firefox for Transparent Authentication

To configure Firefox for NTLM transparent authentication:

Procedure

Step 1 Open **about:config**. Use the filter bar to help you locate the preferences that you need to modify.

Step 2 To support NTLM, modify the following preferences (filter on `network.automatic`):

- **network.automatic-ntlm-auth.trusted-uris**—Double-click the preference, enter the URL, and click **OK**. You can enter multiple URLs by separating them with commas; including the protocol is optional. For example:

```
http://host.example.com, http://hostname, myhost.example.com
```

You can also use partial URLs. Firefox matches the end of the string, not a random substring. Thus, you could include your entire internal network by specifying just your domain name. For example:

```
example.com
```

- **network.automatic-ntlm-auth.allow-proxies**—Ensure that the value is **true**, which is the default. Double-click to change the value if it is currently false.

Step 3 Check the HTTP proxy settings. You can find these by selecting **Tools > Options**, then click the **Network** tab in the Options dialog box. Click the **Settings** button in the Connection group.

- If **No Proxy** is selected, there is nothing to configure.
 - If **Use System Proxy Settings** is selected, you need to modify the **network.proxy.no_proxies_on** property in `about:config` to add the trusted URIs you included in **network.automatic-ntlm-auth.trusted-uris**.
 - If **Manual Proxy Configuration** is selected, update the **No Proxy For** list to include these trusted URIs.
 - If one of the other options is selected, ensure that the properties used for those configurations exclude the same trusted URIs.
-

Monitoring Identity Policies

If identity policies that require authentication are working correctly, you should see user information on the **Monitoring > Users** dashboard and other dashboards that include user information.

In addition, events shown in **Monitoring > Events** should include user information.

If you do not see any user information, verify that the directory server is functioning correctly. Use the **Test** button in the directory server configuration dialog box to verify connectivity.

If the directory server is functioning and usable, verify that the traffic matching criteria on the identity rules that require active authentication are written in a way that will match your users. For example, ensure that the source zone contains the interfaces through which your user traffic will enter the device. The active authentication identity rules match HTTP traffic only, so users must be sending that type of traffic through the device.

Examples for Identity Policies

The use case chapter includes an example of implementing identity policies. Please see [How to Gain Insight Into Your Network Traffic, on page 36](#).



CHAPTER 9

Access Control

The following topics explain access control rules. These rules control which traffic is allowed to pass through the device, and apply advanced services to the traffic, such as intrusion inspection.

- [Access Control Overview, on page 139](#)
- [License Requirements for Access Control, on page 146](#)
- [Guidelines and Limitations for Access Control Policies, on page 147](#)
- [Configuring the Access Control Policy, on page 148](#)
- [Monitoring Access Control Policies, on page 158](#)
- [Examples for Access Control, on page 159](#)

Access Control Overview

The following topics explain access control policies.

Access Control Rules and the Default Action

Use the access control policy to allow or block access to network resources. The policy consists of a set of ordered rules, which are evaluated from top to bottom. The rule applied to traffic is the first one where all the traffic criteria are matched.

You can control access based on:

- Traditional network characteristics such as source and destination IP addresses, protocol, ports, and interfaces (in the form of security zones).
- The application that is being used. You can control access based on the specific application, or you can create rules that cover categories of applications, applications tagged with a particular characteristic, the type of application (client, server, web), or the application's risk or business relevance rating.
- The destination URL of a web request, including the generalized category of the URL. You can refine category matches based on the public reputation of the target site.
- The user who is making the request, or the user groups to which the user belongs.

For unencrypted traffic that you allow, you can apply IPS inspection to check for threats and block traffic that appears to be an attack. You can also use file policies to check for prohibited files or malware.

Any traffic that does not match an access rule is handled by the access control **Default Action**. If you allow traffic by default, you can apply intrusion inspection to the traffic. However, you cannot perform file or malware inspection on traffic handled by the default action.

Application Filtering

You can use access control rules to filter traffic based on the application used in the connection. The system can recognize a wide variety of applications, so that you do not need to figure out how to block one web application without blocking all web applications.

For some popular applications, you can filter on different aspects of the application. For example, you could create a rule that blocks Facebook Games without blocking all of Facebook.

You can also create rules based on general application characteristics, blocking or allowing entire groups of applications by selecting risk or business relevance, type, category, or tag. **However, as you select categories in an application filter, look over the list of matching applications to ensure you are not including unintended applications.** For a detailed explanation of the possible groupings, see [Application Criteria, on page 152](#).

Application Control for Encrypted and Decrypted Traffic

If an application uses encryption, the system might not be able to identify the application.

The system can detect application traffic encrypted with StartTLS, including SMTPS, POPS, FTPS, TelnetS, and IMAPS. In addition, it can identify certain encrypted applications based on the Server Name Indication in the TLS ClientHello message, or the subject distinguished name value from the server certificate.

Use the application filters dialog box to determine if your application requires decryption by selecting the following Tags, then examining the list of applications.

- **SSL Protocol**—You do not need to decrypt traffic tagged as SSL Protocol. The system can recognize this traffic and apply your access control action. Access control rules for the listed applications should match to expected connections.
- **Decrypted Traffic**—The system can recognize this traffic only if you first decrypt the traffic. Because you cannot configure SSL decryption using FDM, access control rules for these applications do not work. For example, at the time of this writing, Dropbox has this tag. Thus, access rules for the Dropbox application will not match Dropbox connections.

Best Practices for Application Filtering

Please keep the following recommendations in mind when designing your application filtering access control rules.

- To handle traffic referred by a web server, such as advertisement traffic, match the referred application rather than the referring application.
- Avoid combining application and URL criteria in the same rule, especially for encrypted traffic.
- If you write a rule for traffic that is tagged **Decrypted Traffic**, ensure that you have an SSL Decryption rule that will decrypt the matching traffic. These applications can be identified in decrypted connections only.

- The system can detect multiple types of Skype application traffic. To control Skype traffic, choose the Skype tag from the Application Filters list rather than selecting individual applications. This ensures that the system can detect and control all Skype traffic the same way.
- To control access to Zoho mail, select both the Zoho and Zoho Mail applications.

URL Filtering

You can use access control rules to filter traffic based on the URL used in an HTTP or HTTPS connection. Note that URL filtering for HTTP is more straight-forward than it is for HTTPS, because HTTPS is encrypted.

You can use the following techniques to implement URL filtering.

- Category and reputation-based URL filtering—With a URL filtering license, you can control access to web sites based on the URL's general classification (category) and risk level (reputation). This is by far the easiest and most effective way to block unwanted sites.
- Manual URL filtering—With any license, you can manually specify individual URLs, and groups of URLs, to achieve granular, custom control over web traffic. The main purpose of manual filtering is to create exceptions to category-based block rules, but you can use manual rules for other purposes.

The following topics provide more information on URL filtering.

Filtering URLs by Category and Reputation

With a URL filtering license, you can control access to web sites based on the category and reputation of the requested URLs:

- Category—A general classification for the URL. For example, ebay.com belongs to the Auctions category, and monster.com belongs to the Job Search category. A URL can belong to more than one category.
- Reputation—How likely the URL is to be used for purposes that might be against your organization's security policy. Reputations range from High Risk (level 1) to Well Known (level 5).

URL categories and reputations help you quickly configure URL filtering. For example, you can use access control to block high risk URLs in the Abused Drugs category.

Using category and reputation data also simplifies policy creation and administration. Sites that represent security threats, or that serve undesirable content, might appear and disappear faster than you can update and deploy new policies. As Cisco updates the URL database with new sites, changed classifications, and changed reputations, your rules automatically adjust to the new information. You do not need to edit your rules to account for new sites.

If you enable regular URL database updates, you can ensure that the system uses up-to-date information for URL filtering. You can also enable communications with Cisco Collective Security Intelligence (CSI) to obtain the latest threat intelligence for URLs with unknown category and reputation. For more information, see [Configuring URL Filtering Preferences, on page 289](#).



Note To see URL category and reputation information in events and application details, you must create at least one rule with a URL condition.

Looking Up the Category and Reputation for a URL

You can check on the category and reputation for a particular URL by using the following site. You can use this information to help you check the behavior of your category and reputation based URL filtering rules.

<https://www.brightcloud.com/tools/url-ip-lookup.php>

Manual URL Filtering

You can supplement or selectively override category and reputation-based URL filtering by manually filtering individual URLs or groups of URLs. You can perform this type of URL filtering without a special license.

For example, you might use access control to block a category of web sites that are not appropriate for your organization. However, if the category contains a web site that is appropriate, and to which you want to provide access, you can create a manual Allow rule for that site and place it before the Block rule for the category.

To configure manual URL filtering, you create a URL object with the destination URL. How this URL is interpreted is based on the following rules:

- If you do not include a path (that is, there is no / character in the URL), the match is based on the server's hostname only. The hostname is considered a match if it comes after the `://` separator, or after any dot in the hostname. For example, `ign.com` matches `ign.com` and `www.ign.com`, but it does not match `verisign.com`.
- If you include one or more / character, the entire URL string is used for a substring match, including the server name, path, and any query parameters. However, we recommend that you do not use manual URL filtering to block or allow individual web pages or parts of sites, as servers can be reorganized and pages moved to new paths. Substring matching can also lead to unexpected matches, where the string you include in the URL object also matches paths on unintended servers or strings within query parameters.
- The system disregards the encryption protocol (HTTP vs HTTPS). In other words, if you block a website, both HTTP and HTTPS traffic to that website is blocked, unless you use an application condition to target a specific protocol. When creating a URL object, you do not need to specify the protocol when creating an object. For example, use `example.com` rather than `http://example.com`.
- If you plan to use a URL object to match HTTPS traffic in an access control rule, create the object using the subject common name in the public key certificate used to encrypt the traffic. Also, the system disregards subdomains within the subject common name, so do not include subdomain information. For example, use `example.com` rather than `www.example.com`.

However, please understand that the subject common name in the certificate might be completely unrelated to a web site's domain name. For example, the subject common name in the certificate for `youtube.com` is `*.google.com` (this of course might change at any time). You will get more consistent results if you use the SSL Decryption policy to decrypt HTTPS traffic so that URL filtering rules work on decrypted traffic.



Note URL objects will not match HTTPS traffic if the browser resumes a TLS session because the certificate information is no longer available. Thus, even if you carefully configure the URL object, you might get inconsistent results for HTTPS connections.

Filtering HTTPS Traffic

Because HTTPS traffic is encrypted, performing URL filtering directly on HTTPS traffic is not as straight-forward as it is on HTTP traffic. For that reason, you should consider using SSL Decryption policies to decrypt all HTTPS traffic that you intend to filter. That way, the URL filtering access control policies work on decrypted traffic, and you get the same results you would get for regular HTTP traffic.

However, if you do intend to allow some HTTPS traffic to pass undecrypted into the access control policy, you need to understand that rules match HTTPS traffic differently than they do for HTTP traffic. To filter encrypted traffic, the system determines the requested URL based on information passed during the SSL handshake: the subject common name in the public key certificate used to encrypt the traffic. There might be little or no relationship between the web site hostname in the URL and the subject common name.

HTTPS filtering, unlike HTTP filtering, disregards subdomains within the subject common name. Do not include subdomain information when manually filtering HTTPS URLs. For example, use `example.com` rather than `www.example.com`. Also, review the content of the certificates used by the site to ensure you have the right domain, the one used in the subject common name, and that this name will not conflict with your other rules (for example, the name for a site you want to block might overlap with one you want to allow). For example, the subject common name in the certificate for `youtube.com` is `*.google.com` (this of course might change at any time).



Note URL objects will not match HTTPS traffic if the browser resumes a TLS session because the certificate information is no longer available. Thus, even if you carefully configure the URL object, you might get inconsistent results for HTTPS connections.

Controlling Traffic by Encryption Protocol

The system disregards the encryption protocol (HTTP vs HTTPS) when performing URL filtering. This occurs for both manual and reputation-based URL conditions. In other words, URL filtering treats traffic to the following web sites identically:

- `http://example.com`
- `https://example.com`

To configure a rule that matches only HTTP or HTTPS traffic, but not both, either specify the TCP port in the Destination condition or add an application condition to the rule. For example, you could allow HTTPS access to a site while disallowing HTTP access by constructing two access control rules, each with an TCP port or application, and URL, condition.

The first rule allows HTTPS traffic to the website:

Action: Allow
TCP port or Application: HTTPS (TCP port 443)
URL: `example.com`

The second rule blocks HTTP access to the same website:

Action: Block
TCP port or Application: HTTP (TCP port 80)
URL: `example.com`

Comparing URL and Application Filtering

URL and application filtering have similarities. But you should use them for very distinct purposes:

- URL filtering is best used to block or allow access to an entire web server. For example, if you do not want to allow any type of gambling on your network, you can create a URL filtering rule to block the Gambling category. With this rule, users cannot get to any pages on any web server within the category.
- Application filtering is useful for blocking specific applications regardless of the hosting site, or for blocking specific features of an otherwise allowable web site. For example, you could block just the Facebook Games application without blocking all of Facebook.

Because combining application and URL criteria can lead to unexpected results, especially for encrypted traffic, it is a good policy to create separate rules for URL and application criteria. If you do need to combine application and URL criteria in a single rule, you should place these rules after straight-forward application-only or URL-only rules, unless the application+URL rule is acting as an exception to a more general application-only or URL-only rule. Because URL filtering block rules are more broad than application filtering, you should place them above application-only rules.

If you do combine application and URL criteria, you might need to monitor your network more carefully to ensure that you are not allowing access to unwanted sites and applications.

Best Practices for Effective URL Filtering

Please keep the following recommendations in mind when designing your URL filtering access control rules.

- Use category and reputation blocking whenever possible. This ensures that new sites get blocked automatically as they are added to the categories, and that blocking based on reputation is adjusted if a site becomes more (or less) reputable.
- When using URL category matching, note that there are cases where the login page for a site is in a different category than the site itself. For example, Gmail is in the Web-based Email category, whereas the login page is in the Internet Portals category. If you have different rules with different actions for the categories, you might get unintended results.
- Use URL objects to target entire web sites and to make exceptions to category blocking rules. That is, to allow specific sites that would otherwise get blocked in a category rule.
- For the most effective filtering of HTTPS connections, implement SSL decryption rules to decrypt traffic for which you are writing an access control rule. Any decrypted HTTPS connections are filtered as HTTP connections in the access control policy, so you avoid all of the limitations for HTTPS filtering.
- Place URL blocking rules before any application filtering rules, because URL filtering blocks entire web servers, whereas application filtering targets specific application usage regardless of the web server.

What the User Sees When You Block Web Sites

When you block web sites with URL filtering rules, what the user sees differs based on whether the site is encrypted.

- HTTP connections—The user sees a system default block response page instead of the normal browser page for timed out or reset connections. This page should make it clear that you blocked the connection on purpose.
- HTTPS (encrypted) connections—The user does not see the system default block response page. Instead, the user sees the browser's default page for a secure connection failure. The error message does not

indicate the site was blocked due to policy. Instead, errors might indicate that there are no common encryption algorithms. It will not be obvious from this message that you blocked the connection on purpose.

In addition, web sites might be blocked by other access control rules that are not explicitly URL filtering rules, or even by the default action. For example, if you block entire networks or geolocations, any web sites on that network or in that geographic location are also blocked. Users blocked by these rules may, or may not, get a response page as described in the limitations below.

If you implement URL filtering, consider explaining to end users what they might see when a site is intentionally blocked, and what types of site you are blocking. Otherwise, they might spend a good deal of time troubleshooting blocked connections.

Limitations of HTTP Response Pages

HTTP response pages do not always appear when the system blocks web traffic.

- The system does not display a response page when web traffic is blocked as a result of a promoted access control rule (an early-placed blocking rule with only simple network conditions).
- The system does not display a response page when web traffic is blocked before the system identifies the requested URL.
- The system does not display a response page for encrypted connections blocked by access control rules.

Intrusion, File, and Malware Inspection

Intrusion and file policies work together as the last line of defense before traffic is allowed to its destination:

- Intrusion policies govern the system's intrusion prevention capabilities.
- File policies govern the system's file control and malware defense capabilities.

All other traffic handling occurs before network traffic is examined for intrusions, prohibited files, and malware. By associating an intrusion or file policy with an access control rule, you are telling the system that before it passes traffic that matches the access control rule's conditions, you first want to inspect the traffic with an intrusion policy, a file policy, or both.

You can configure intrusion and file policies on rules that **allow** traffic only. Inspection is not performed on rules set to **trust** or **block** traffic. In addition, if the default action for the access control policy is **allow**, you can configure an intrusion policy but not a file policy.

For any single connection handled by an access control rule, file inspection occurs before intrusion inspection. That is, the system does not inspect files blocked by a file policy for intrusions. Within file inspection, simple blocking by type takes precedence over malware inspection and blocking. Until a file is detected and blocked in a session, packets from the session may be subject to intrusion inspection.



Note By default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured. Inspection works with unencrypted traffic only.

Best Practices for Access Control Rule Order

Rules are applied on a first-match basis, so you must ensure that rules with highly specific traffic matching criteria appear above policies that have more general criteria that would otherwise apply to the matching traffic. Consider the following recommendations:

- Specific rules should come before general rules, especially when the specific rules are exceptions to general rules.
- Any rules that drop traffic based on layer-3/4 criteria only (such as IP address, security zone, and port number) should come as early as possible. We recommend they come before any rule that requires inspection, such as those with application or URL criteria, because Layer-3/4 criteria can be evaluated quickly and without inspection. Of course, any exceptions to these rules must be placed above them.
- Whenever possible, put specific drop rules near the top of the policy. This ensures the earliest possible decision on undesirable traffic.
- Any rules that include both application and URL criteria should come after straight-forward application-only or URL-only rules, unless the application+URL rule is acting as an exception to a more general application-only or URL-only rule. Combining application and URL criteria can lead to unexpected results, especially for encrypted traffic, so we recommend that you create separate rules for URL and application filtering whenever possible.

NAT and Access Rules

Access rules always use the real IP addresses when determining an access rule match, even if you configure NAT. For example, if you configure NAT for an inside server, 10.1.1.5, so that it has a publicly routable IP address on the outside, 209.165.201.5, then the access rule to allow the outside traffic to access the inside server needs to reference the server's real IP address (10.1.1.5), and not the mapped address (209.165.201.5).

How Other Security Policies Impact Access Control

Other security policies can affect how access control rules function and match connections. As you configure your access rules, keep the following in mind:

- **Identity** policy—Connections are matched to users (and thus, user groups) only if there is a user mapping for the source IP address. Access rules that key on user or group membership can match only those connections for which user identity was successfully collected by your identity policy.
- **VPN (site-to-site)**—VPN traffic is always evaluated against the access control policy, and connections are allowed or dropped based on the matching rule. However, the VPN tunnel itself is decrypted before the access control policy is evaluated. The access control policy evaluates the connections that are embedded within the VPN tunnel, not the tunnel itself.

License Requirements for Access Control

You do not need a special license to use the access control policy.

However, you do need the following licenses for specific features within the access control policy. For information on configuring licenses, see [Enabling or Disabling Optional Licenses, on page 67](#).

- **URL** license—To create rules that use URL categories and reputations as match criteria.
- **Threat** license—To configure an intrusion policy on an access rule or the default action. You also need this license to use a file policy (the Malware license is also required).
- **Malware** license—To configure a file policy on an access rule. The Threat is also required for file policies.

Guidelines and Limitations for Access Control Policies

Following are some additional limitations for access control. Please consider them when evaluating whether you are getting the expected results from your rules.

- FDM can download information on up to 2000 users from the directory server. If your directory server includes more than 2000 user accounts, you will not see all possible names when selecting users in an access rule or when viewing user-based dashboard information. You can write rules on only those names that were downloaded.

The 2000 limit also applies to the names associated with groups. If a group has more than 2000 members, only the 2000 names that were downloaded can be matched against the group membership.

- If a Vulnerability Database (VDB) update removes (deprecates) applications, you must make changes to any access control rules or application filters that use the application that was deleted. You cannot deploy changes until you fix these rules. In addition, you cannot install system software updates before fixing the issue. On the Application Filters object page, or the Application tab of the rule, these applications say “(Deprecated)” after the application name.
- If you edit a rule that is actively in use, the changes do not apply to established connections that are no longer being inspected by Snort. The new rule is used to match against future connections. In addition, if Snort is actively inspecting a connection, it can apply the changed matching or action criteria to an existing connection. If you need to ensure that your changes apply to all current connections, you can log into the device CLI and use the **clear conn** command to end established connections, on the assumption that the sources for the connections will then attempt to reestablish the connection and thus be matched appropriately against the new rule.
- It takes 3 to 5 packets for the system to identify the application or URL in a connection. Thus, the correct access control rule might not be matched immediately for a given connection. However, once the application/URL is known, the connection is handled based on the matching rule. For encrypted connections, this happens after the server certificate exchange in the SSL handshake.
- The system applies the default policy action to packets that do not have a payload in a connection where an application is identified.
- Leave matching criteria empty whenever possible, especially those for security zones, network objects, and port objects. For example, the system can more efficiently match traffic for all interfaces if you simply leave the security zone criteria blank, rather than if you create zones that contain all interfaces. When you specify multiple criteria, the system must match against every combination of the contents of the criteria you specify.
- If you specify IP addresses for source or destination criteria, do not mix IPv4 and IPv6 addresses in the same rule. Create separate rules for IPv4 and IPv6 addresses.
- Due to memory limitations, some device models perform most URL filtering with a smaller, less granular, set of categories and reputations. For example, even if a parent URL's subsites have different URL

categories and reputations, some devices may only store the parent URL's data. For web traffic handled by these devices, the system may perform cloud lookups to determine category and reputation for sites not in the local database. Lower-memory devices include the following ASA models: 5506-X, 5506H-X, 5506W-X, 5508-X, 5512-X, 5515-X, 5516-X, and 5525-X.




- GRE tunnels that violate the related RFCs will be dropped. For example, if a GRE tunnel contains non-zero values in the reserved bits, contrary to the RFCs, it is dropped. If you need to allow non-compliant GRE tunnels, you need to use a remote manager and configure a prefilter rule that trusts the sessions. You cannot configure prefilter rules using the FDM.

Configuring the Access Control Policy

Use the access control policy to control access to network resources. The policy consists of a set of ordered rules, which are evaluated from top to bottom. The rule applied to traffic is the first one where all the traffic criteria are matched. If no rules match the traffic, the default action shown at the bottom of the page is applied.

To configure the access control policy, select **Policies > Access Control**.

The access control table lists all rules in order. For each rule:

- Click the > button next to the rule number in the left-most column to open the rule diagram. The diagram can help you visualize how the rule controls traffic. Click the button again to close the diagram.
- Most cells allow inline editing. For example, you can click the action to select a different one, or click a source network object to add or change the source criteria.
- To move a rule, hover over the rule until you get the move icon () , then click, drag, and drop the rule to the new location. You can also move a rule by editing it and selecting the new location in the **Order** list. It is critical that you put the rules in the order that you want them processed. Specific rules should be near the top, especially for rules that define exceptions to more general rules
- The right-most column contains the action buttons for a rule; mouse over the cell to see the buttons. You can edit () or delete () a rule.

The following topics explain how to configure the policy.

Configuring the Default Action

If a connection does not match a specific access rule, it is handled by the default action for the access control policy.

Procedure

-
- Step 1** Select **Policies > Access Control**.
 - Step 2** Click anywhere in the **Default Action** field.
 - Step 3** Select the action to apply to matching traffic.
 - **Trust**—Allow traffic without further inspection of any kind.
 - **Allow**—Allow the traffic subject to the intrusion policy.



- **Block**—Drop the traffic unconditionally. The traffic is not inspected.

- Step 4** If the action is **Allow**, select **Enable Policy > On** under **Intrusion Policy** and select an intrusion policy. For an explanation of the policy options, see [Intrusion Policy Settings, on page 155](#).
- Step 5** (Optional.) Configure logging for the default action. You must enable logging for traffic that matches the default action to be included in dashboard data or Event Viewer. See [Logging Settings, on page 157](#).
- Step 6** Click **OK**.
-

Configuring Access Control Rules

Use access control rules to control access to network resources. Rules in the access control policy are evaluated from top to bottom. The rule applied to traffic is the first one where all the traffic criteria are matched.

Procedure

- Step 1** Select **Policies > Access Control**.
- Step 2** Do any of the following:
- To create a new rule, click the + button.
 - To edit an existing rule, click the edit icon () for the rule.
- To delete a rule you no longer need, click the delete icon () for the rule.
- Step 3** In **Order**, select where you want to insert the rule in the ordered list of rules. Rules are applied on a first-match basis, so you must ensure that rules with highly specific traffic matching criteria appear above policies that have more general criteria that would otherwise apply to the matching traffic. The default is to add the rule to the end of the list. If you want to change a rule's location later, edit this option.
- Step 4** In **Title**, enter a name for the rule. The name cannot contain spaces. You can use alphanumeric characters and these special characters: + . _ -
- Step 5** Select the action to apply to matching traffic.
- **Trust**—Allow traffic without further inspection of any kind.
 - **Allow**—Allow the traffic subject to the intrusion and other inspection settings in the policy.
 - **Block**—Drop the traffic unconditionally. The traffic is not inspected.
- Step 6** Define the traffic matching criteria using any combination of the following tabs:
- **Source/Destination**—The security zones (interfaces) through which the traffic passes, the IP addresses or the country or continent (geographical location) for the IP address, or the protocols and ports used in the traffic. The default is any zone, address, geographical location, protocol, and port. See [Source/Destination Criteria, on page 150](#).

- **Application**—The application, or a filter that defines applications by type, category, tag, risk, or business relevance. The default is any application. See [Application Criteria, on page 152](#).
- **URL**—The URL or URL category of a web request. The default is any URL. See [URL Criteria, on page 153](#).
- **Users**—The user or user group. Your identity policies determine whether user and group information is available for traffic matching. You must configure identity policies to use this criteria. See [User Criteria, on page 154](#).

To modify a condition, you click the + button within that condition, select the desired object or element, and click **OK** in the popup dialog box. If the criterion requires an object, you can click **Create New Object** if the object you require does not exist. Click the x for an object or element to remove it from the policy.

When adding conditions to access control rules, consider the following tips:

- You can configure multiple conditions per rule. Traffic must match all the conditions in the rule for the rule to apply to traffic. For example, you can use a single rule to perform URL filtering for specific hosts or networks.
- For each condition in a rule, you can add up to 50 criteria. Traffic that matches any of a condition's criteria satisfies the condition. For example, you can use a single rule to apply application control for up to 50 applications or application filters. Thus, there is an OR relationship among the items in a single condition, but an AND relationship between condition types (for example, between source/destination and application).
- Some features require that you enable the appropriate license.

Step 7 (Optional.) For policies that use the Allow action, you can configure further inspection on unencrypted traffic. Click one of the following links:

- **Intrusion Policy**—Select **Intrusion Policy > On** and select the intrusion inspection policy to inspect traffic for intrusions and exploits. See [Intrusion Policy Settings, on page 155](#).
- **File Policy**—Select the file policy to inspect traffic for files that contain malware and for files that should be blocked. See [File Policy Settings, on page 156](#).

Step 8 (Optional.) Configure logging for the rule.

By default, connection events are not generated for traffic that matches a rule, although file events are generated by default if you select a file policy. You can change this behavior. You must enable logging for traffic that matches the policy to be included in dashboard data or Event Viewer. See [Logging Settings, on page 157](#).

Intrusion events are always generated for intrusion rules set to drop or alert regardless of the logging configuration on the matching access rule.

Step 9 Click **OK**.

Source/Destination Criteria

The Source/Destination criteria of an access rule define the security zones (interfaces) through which the traffic passes, the IP addresses or the country or continent (geographical location) for the IP address, or the protocols and ports used in the traffic. The default is any zone, address, geographical location, protocol, and port.

To modify a condition, you click the + button within that condition, select the desired object or element, and click **OK**. If the criterion requires an object, you can click **Create New Object** if the object you require does not exist. Click the **x** for an object or element to remove it from the policy.

You can use the following criteria to identify the source and destination to match in the rule.

Source Zones, Destination Zones

The security zone objects that define the interfaces through which the traffic passes. You can define one, both, or neither criteria: any criteria not specified applies to traffic on any interface.

- To match traffic leaving the device from an interface in the zone, add that zone to the **Destination Zones**.
- To match traffic entering the device from an interface in the zone, add that zone to the **Source Zones**.
- If you add both source and destination zone conditions to a rule, matching traffic must originate from one of the specified source zones and egress through one of the destination zones.

Use this criteria when the rule should apply based on where the traffic enters or exits the device. For example, if you want to ensure that all traffic going to inside hosts gets intrusion inspection, you would select your inside zone as the **Destination Zones** while leaving the source zone empty. To implement intrusion filtering in the rule, the rule action must be **Allow**, and you must select an intrusion policy in the rule.

Source Networks, Destination Networks

The network objects or geographical locations that define the network addresses or locations of the traffic.

- To match traffic from an IP address or geographical location, configure the **Source Networks**.
- To match traffic to an IP address or geographical location, configure the **Destination Networks**.
- If you add both source and destination network conditions to a rule, matching traffic must originate from one of the specified IP addresses and be destined for one of the destination IP addresses.

When you add this criteria, you select from the following tabs:

- **Network**—Select the network objects or groups that define the source or destination IP addresses for the traffic you want to control.
- **Geolocation**—Select the geographical location to control traffic based on its source or destination country or continent. Selecting a continent selects all countries within the continent. Besides selecting geographical location directly in the rule, you can also select a geolocation object that you created to define the location. Using geographical location, you could easily restrict access to a particular country without needing to know all of the potential IP addresses used there.



Note To ensure that you are using up-to-date geographical location data to filter your traffic, Cisco strongly recommends that you regularly update the geolocation database (GeoDB).

Source Ports, Destination Ports/Protocols

The port objects that define the protocols used in the traffic. For TCP/UDP, this can include ports. For ICMP, it can include codes and types.

- To match traffic from a protocol or port, configure the **Source Ports**. Source ports can be TCP/UDP only.
- To match traffic to a protocol or port, configure the **Destination Ports/Protocols**. If you add only destination ports to a condition, you can add ports that use different transport protocols. ICMP and other non-TCP/UDP specifications are allowed in destination ports only; they are not allowed in source ports.
- To match traffic both originating from specific TCP/UDP ports and destined for specific TCP/UDP ports, configure both. If you add both source and destination ports to a condition, you can only add ports that share a single transport protocol, TCP or UDP. For example, you could target traffic from port TCP/80 to port TCP/8080.

Application Criteria

The Application criteria of an access rule defines the application used in an IP connection, or a filter that defines applications by type, category, tag, risk, or business relevance. The default is any application.

Although you can specify individual applications in the rule, application filters simplify policy creation and administration. For example, you could create an access control rule that identifies and blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the session is blocked.

In addition, Cisco frequently updates and adds additional application detectors via system and vulnerability database (VDB) updates. Thus, a rule blocking high risk applications can automatically apply to new applications without you having to update the rule manually.

You can specify applications and filters directly in the rule, or create application filter objects that define those characteristics. The specifications are equivalent, although using objects can make it easier to stay within the 50-items-per-criteria system limit if you are creating a complex rule.

To modify the application and filters list, you click the + button within the condition, select the desired applications or application filter objects, which are listed on separate tabs, and click **OK** in the popup dialog box. On either tab, you can click **Advanced Filter** to select filter criteria or to help you search for specific applications. Click the **x** for an application, filter, or object to remove it from the policy. Click the **Save As Filter** link to save the combined criteria that is not already an object as a new application filter object.



Note If a selected application was removed by a VDB update, “(Deprecated)” appears after the application name. You must remove these applications from the filter, or subsequent deployments and system software upgrades will be blocked.

You can use the following **Advanced Filter** criteria to identify the application or filter to match in the rule. These are the same elements used in application filter objects.



Note Multiple selections within a single filter criteria have an OR relationship. For example, Risk is High OR Very High. The relationship between filters is AND, so Risk is High OR Very High, AND Business Relevance is Low OR Very Low. As you select filters, the list of applications in the display updates to show only those that meet the criteria. You can use these filters to help you find applications that you want to add individually, or to verify that you are selecting the desired filters to add to the rule.

Risks

The likelihood that the application is used for purposes that might be against your organization's security policy, from very low to very high.

Business Relevance

The likelihood that the application is used within the context of your organization's business operations, as opposed to recreationally, from very low to very high.

Types

The type of application:

- **Application Protocol**—Application protocols such as HTTP and SSH, which represent communications between hosts.
- **Client Protocol**—Clients such as web browsers and email clients, which represent software running on the host.
- **Web Application**—Web applications such as MPEG video and Facebook, which represent the content or requested URL for HTTP traffic.

Categories

A general classification for the application that describes its most essential function.

Tags

Additional information about the application, similar to category.

For encrypted traffic, the system can identify and filter traffic using only the applications tagged **SSL Protocol**. Applications without this tag can only be detected in unencrypted or decrypted traffic. Also, the system assigns the **decrypted traffic** tag to applications that the system can detect in decrypted traffic only, not encrypted or unencrypted.

Applications List (bottom of the display)

This list updates as you select filters from the options above the list, so you can see the applications that currently match the filter. Use this list to verify that your filter is targeting the desired applications when you intend to add filter criteria to the rule. If your intention is to add specific applications, select them from this list.

URL Criteria

The URL criteria of an access rule defines the URL used in a web request, or the category to which the requested URL belongs. For category matches, you can also specify the relative reputation of sites to allow or block. The default is to allow all URLs.

URL categories and reputations allow you to quickly create URL conditions for access control rules. For example, you could block all Gambling sites, or high risk Social Networking sites. If a user attempts to browse to any URL with that category and reputation combination, the session is blocked.

Using category and reputation data also simplifies policy creation and administration. It grants you assurance that the system will control web traffic as expected. Finally, because Cisco's threat intelligence is continually updated with new URLs, as well as new categories and risks for existing URLs, you can ensure that the system uses up-to-date information to filter requested URLs. Malicious sites that represent security threats such as malware, spam, botnets, and phishing may appear and disappear faster than you can update and deploy new policies.

To modify the URL list, you click the + button within the condition and select the desired categories or URLs using one of the following techniques. Click the **x** for a category or object to remove it from the policy.

URL Tab

Click +, select URL objects or groups, and click **OK**. You can click **Create New URL** if the object you require does not exist.



Note Before configuring URL objects to target specific sites, carefully read the information on manual URL filtering.

Categories Tab

Click +, select the desired categories, and click **OK**.

The default is to apply the rule to all URLs in each selected category regardless of reputation. To limit the rule based on reputation, click the down arrow for each category, deselect the **Any** checkbox, and then use the **Reputation** slider to choose the reputation level. The left of the reputation slider indicates sites that will be allowed, the right side are sites that will be blocked. How reputation is used depends on the rule action:

- If the rule blocks or monitors web access, selecting a reputation level also selects all reputations more severe than that level. For example, if you configure a rule to block or monitor **Suspicious sites** (level 2), it also automatically blocks or monitors **High risk** (level 1) sites.
- If the rule allows web access, selecting a reputation level also selects all reputations less severe than that level. For example, if you configure a rule to allow **Benign sites** (level 4), it also automatically allows **Well known** (level 5) sites.

User Criteria

The User criteria of an access rule defines the user or user group for an IP connection. You must configure identity policies and the associated directory server to include user or user group criteria in an access rule.

Your identity policies determine whether user identity is collected for a particular connection. If identity is established, the IP address of the host is associated with the identified user. Thus, traffic whose source IP address is mapped to a user is considered to be from that user. IP packets themselves do not include user identity information, so this IP-address-to-user mapping is the best approximation available.

Because you can add a maximum of 50 users or groups to a rule, selecting groups usually makes more sense than selecting individual users. For example, you could create a rule allowing the Engineering group access to a development network, and create a subsequent rule that denies all other access to the network. Then, to

make the rule apply to new engineers, you only need to add the engineer to the Engineering group in the directory server.

To modify the users list, you click the + button within the condition and select the desired identities using one of the following techniques. Click the x for an identity to remove it from the policy.

- **Users and Groups** tab—Select the desired users or user groups. Groups are available only if you configure groups in the directory server. If you select a group, the rule applies to any member of the group, including subgroups. If you want to treat a sub-group differently, you need to create a separate access rule for the sub-group and place it above the rule for the parent group in the access control policy.
- **Special Entities** tab—Select from the following:
 - **Failed Authentication**—The user was prompted to authenticate, but failed to enter a valid username/password pair within the maximum number of allowed attempts. Failure to authenticate does not itself prevent the user from accessing the network, but you can write an access rule to limit network access for these users.
 - **Guest**—Guest users are like Failed Authentication users, except that your identity rule is configured to call these users Guest. Guest users were prompted to authenticate and failed to do so within the maximum number of attempts.
 - **No Authentication Required**—The user was not prompted to authentication, because the user's connections matched identity rules that specified no authentication.
 - **Unknown**—There is no user mapping for the IP address, and there is no record of failed authentication yet. Typically, this means that no HTTP traffic has yet been seen from that address.

Intrusion Policy Settings

Cisco delivers several intrusion policies with the Firepower System. These policies are designed by the Cisco Talos Intelligence Group (Talos), who set the intrusion and preprocessor rule states and advanced settings. You cannot modify these policies.

For access control rules that allow traffic, you can select one of the following intrusion policies to inspect traffic for intrusions and exploits. An intrusion policy examines decoded packets for attacks based on patterns, and can block or alter malicious traffic.

To enable intrusion inspection, select **Intrusion Policy > On** and select the desired policy. The policies are listed from least to most secure.

- **Connectivity over Security**—This policy is built for organizations where connectivity (being able to get to all resources) takes precedence over network infrastructure security. The intrusion policy enables far fewer rules than those enabled in the Security over Connectivity policy. Only the most critical rules that block traffic are enabled. Select this policy if you want to apply some intrusion protection but you are fairly confident in the security of your network.
- **Balanced Security and Connectivity**—This policy is designed to balance overall network performance with network infrastructure security. This policy is appropriate for most networks. Select this policy for most situations where you want to apply intrusion prevention.
- **Security over Connectivity**—This policy is built for organizations where network infrastructure security takes precedence over user convenience. The intrusion policy enables numerous network anomaly intrusion rules that could alert on or drop legitimate traffic. Select this policy when security is paramount or for traffic that is high risk.

- **Maximum Detection**—This policy is built for organizations where network infrastructure security is given even more emphasis than is given by the Security Over Connectivity policy, with the potential for even greater operational impact. For example, the intrusion policy enables rules in a large number of threat categories including malware, exploit kit, old and common vulnerabilities, and known in-the-wild exploits. If you select this policy, carefully evaluate whether too much legitimate traffic is being dropped.

File Policy Settings

Use file policies to detect malicious software, or *malware*, using malware defense. You can also use file policies to perform file control, which allows control over all files of a specific type regardless of whether the files contain malware.

Malware defense uses the AMP Cloud to retrieve dispositions for possible malware detected in network traffic, and to obtain local malware analysis and file pre-classification updates. The management interface must have a path to the Internet to reach the AMP Cloud and perform malware lookups. When the device detects an eligible file, it uses the file's SHA-256 hash value to query the AMP Cloud for the file's disposition. The possible dispositions are:

- **Malware**—The AMP Cloud categorized the file as malware. An archive file (e.g. a zip file) is marked as malware if any file within it is malware.
- **Clean**—The AMP Cloud categorized the file as clean, containing no malware. An archive file is marked as clean if all files within it are clean.
- **Unknown**—The AMP Cloud has not assigned a disposition to the file yet. An archive file is marked as unknown if any file within it is unknown.
- **Unavailable**—The system could not query the AMP Cloud to determine the file's disposition. You may see a small percentage of events with this disposition; this is expected behavior. If you see a number of "unavailable" events in succession, ensure that the Internet connection for the management address is functioning correctly.

Available File Policies

You can select one of the following file policies:

- **None**—Do not evaluate transmitted files for malware and do no file-specific blocking. Select this option for rules where file transmissions are trusted or where they are unlikely (or impossible), or for rules where you are confident your application or URL filtering adequately protects your network.
- **Block Malware All**—Query the AMP Cloud to determine if files traversing your network contain malware, then block files that represent threats.
- **Cloud Lookup All**—Query the AMP Cloud to obtain and log the disposition of files traversing your network while still allowing their transmission.
- **Block Office Document and PDF Upload, Block Malware Others**—Block users from uploading Microsoft Office documents and PDFs. Additionally, query the AMP Cloud to determine if files traversing your network contain malware, then block files that represent threats.
- **Block Office Documents Upload, Block Malware Others**—Block users from uploading Microsoft Office documents. Additionally, query the AMP Cloud to determine if files traversing your network contain malware, then block files that represent threats.

Logging Settings

The logging settings for an access rule determine whether connection events are issued for traffic that matches the rule. You must enable logging to see events related to the rule in the Event Viewer. You must also enable logging for matching traffic to be reflected in the various dashboards you can use to monitor the system.

You should log connections according to the security and compliance needs of your organization. If your goal is to limit the number of events you generate and improve performance, only enable logging for the connections critical to your analysis. However, if you want a broad view of your network traffic for profiling purposes, you can enable logging for additional connections.



Caution Logging blocked TCP connections during a Denial of Service (DoS) attack can affect system performance and overwhelm the database with multiple similar events. Before you enable logging for a Block rule, consider whether the rule is for an Internet-facing interface or other interface vulnerable to DoS attack.

You can configure the following logging actions.

Select Log Action

You can select one of the following actions:

- **Log at Beginning and End of Connection**—Issue events at the start and end of a connection. Because end-of-connection events contain everything that start-of-connection events contain, plus all of the information that could be gleaned during the connection, Cisco recommends that you do not select this option for traffic that you are allowing. Logging both events can impact system performance. However, this is the only option allowed for blocked traffic.
- **Log at End of Connection**—Select this option if you want to enable connection logging at the end of the connection, which is recommended for allowed or trusted traffic.
- **No Logging at Connection**—Select this option to disable logging for the rule. This is the default.



Note When an intrusion policy invoked by an access control rule detects an intrusion and generates an intrusion event, the system automatically logs the end of the connection where the intrusion occurred, regardless of the logging configuration of the rule. For connections where an intrusion was blocked, the action for the connection in the connection log is **Block**, with a reason of **Intrusion Block**, even though to perform intrusion inspection you must use an Allow rule.

File Events

Select **Log Files** if you want to enable logging of prohibited files or malware events. You must select a file policy in the rule to configure this option. The option is enabled by default if you select a file policy for the rule. Cisco recommends you leave this option enabled.

When the system detects a prohibited file, it automatically logs one of the following types of event:

- *File events*, which represent detected or blocked files, including malware files.
- *Malware events*, which represent detected or blocked malware files only.
- *Retrospective malware events*, which are generated when the malware disposition for a previously detected file changes.

For connections where a file was blocked, the action for the connection in the connection log is **Block** even though to perform file and malware inspection you must use an Allow rule. The connection's Reason is either **File Monitor** (a file type or malware was detected), or **Malware Block** or **File Block** (a file was blocked).

Send Connection Events To

If you want to send a copy of the events to an external syslog server, select the server object that defines the syslog server. If the required object does not already exist, click **Create New Syslog Server** and create it. (To disable logging to a syslog server, select **Any** from the server list.)

Because event storage on the device is limited, sending events to an external syslog server can provide more long term storage and enhance your event analysis.

Monitoring Access Control Policies

The following topics explain how you can monitor the access control policy.

Monitoring Access Control Statistics in the Dashboards

Most of the data on the **Monitoring** dashboards are directly related to your access control policy. See [Monitoring Traffic and System Dashboards](#), on page 71.

- **Monitoring > Policies** shows the most-hit access control rules and related statistics.
- You can find general statistics on the **Network Overview**, **Destinations**, **Ingress Zones**, and **Egress Zones** dashboards.
- You can find URL filtering results on the **Web Categories** and **Destinations** dashboards. You must have at least one URL filtering policy to see any information on the **Web Categories** dashboard.
- You can find application filtering results on the **Applications** dashboard.
- You can find user-based statistics on the **Users** dashboard. You must implement identity policies to collect user information.
- You can find intrusion policy statistics on the **Attackers** and **Targets** dashboards. You must apply an intrusion policy to at least one access control rule to see any information on these dashboards.
- You can find file policy and malware filtering statistics on the **File Logs** dashboard. You must apply a file policy to at least one access control rule to see any information on this dashboard.
- **Monitoring > Events** also shows events for connections and data related to the access control rules.

Monitoring Access Control Policies in the CLI

You can also log into the device CLI and use the following commands to get more detailed information about access control policies and statistics.

- **show access-control-config** displays summary information about the access control rules along with per-rule hit counts.
- **show access-list** displays the access control lists (ACLs) that were generated from the access control rules. The ACLs provide an initial filter and attempt to provide quick decisions whenever possible, so

that connections that should be dropped do not need to be inspected (and thus consume resources unnecessarily). This information includes hit counts.

- **show snort statistics** displays information about the Snort inspection engine, which is the main inspector. Snort implements application filtering, URL filtering, intrusion protection, and file and malware filtering.
- **show conn** displays information about the connections currently established through the interfaces.
- **show traffic** displays statistics about traffic flowing through each interface.
- **show ipv6 traffic** displays statistics about IPv6 traffic flowing through the device.

Examples for Access Control

The use case chapter includes several examples of implementing access control rules. Please see the following examples:

- [How to Gain Insight Into Your Network Traffic, on page 36](#). This example shows some basic ideas for collecting overall connection and user information.
- [How to Block Threats, on page 42](#). This example shows how to apply intrusion policies.
- [How to Block Malware, on page 46](#). This example shows how to apply file policies.
- [How to Implement an Acceptable Use Policy \(URL Filtering\), on page 49](#). This example shows how to perform URL filtering.
- [How to Control Application Usage, on page 53](#). This example shows how to perform application filtering.
- [How to Add a Subnet, on page 57](#). This example shows how to integrate a new subnet into your overall network, including the access rules needed to allow traffic flow.



CHAPTER 10

Network Address Translation (NAT)

The following topics explain Network Address Translation (NAT) and how to configure it.

- [Why Use NAT?, on page 161](#)
- [NAT Basics, on page 162](#)
- [Guidelines for NAT, on page 168](#)
- [Configure NAT, on page 173](#)
- [Translating IPv6 Networks, on page 198](#)
- [Monitoring NAT, on page 212](#)
- [Examples for NAT, on page 213](#)

Why Use NAT?

Each computer and device within an IP network is assigned a unique IP address that identifies the host. Because of a shortage of public IPv4 addresses, most of these IP addresses are private, not routable anywhere outside of the private company network. RFC 1918 defines the private IP addresses you can use internally that should not be advertised:

- 10.0.0.0 through 10.255.255.255
- 172.16.0.0 through 172.31.255.255
- 192.168.0.0 through 192.168.255.255

One of the main functions of NAT is to enable private IP networks to connect to the Internet. NAT replaces a private IP address with a public IP address, translating the private addresses in the internal private network into legal, routable addresses that can be used on the public Internet. In this way, NAT conserves public addresses because it can be configured to advertise at a minimum only one public address for the entire network to the outside world.

Other functions of NAT include:

- **Security**—Keeping internal IP addresses hidden discourages direct attacks.
- **IP routing solutions**—Overlapping IP addresses are not a problem when you use NAT.
- **Flexibility**—You can change internal IP addressing schemes without affecting the public addresses available externally; for example, for a server accessible to the Internet, you can maintain a fixed IP address for Internet use, but internally, you can change the server address.

- Translating between IPv4 and IPv6 (Routed mode only) —If you want to connect an IPv6 network to an IPv4 network, NAT lets you translate between the two types of addresses.



Note NAT is not required. If you do not configure NAT for a given set of traffic, that traffic will not be translated, but will have all of the security policies applied as normal.

NAT Basics

The following topics explain some of the basics of NAT.

NAT Terminology

This document uses the following terminology:

- Real address/host/network/interface—The real address is the address that is defined on the host, before it is translated. In a typical NAT scenario where you want to translate the inside network when it accesses the outside, the inside network would be the “real” network. Note that you can translate any network connected to the device, not just an inside network. Therefore if you configure NAT to translate outside addresses, “real” can refer to the outside network when it accesses the inside network.
- Mapped address/host/network/interface—The mapped address is the address that the real address is translated to. In a typical NAT scenario where you want to translate the inside network when it accesses the outside, the outside network would be the “mapped” network.



Note During address translation, IP addresses configured for the device interfaces are not translated.

- Bidirectional initiation—Static NAT allows connections to be initiated *bidirectionally*, meaning both to the host and from the host.
- Source and destination NAT—For any given packet, both the source and destination IP addresses are compared to the NAT rules, and one or both can be translated/untranslated. For static NAT, the rule is bidirectional, so be aware that “source” and “destination” are used in commands and descriptions throughout this guide even though a given connection might originate at the “destination” address.

NAT Types

You can implement NAT using the following methods:


- Dynamic NAT—A group of real IP addresses are mapped to a (usually smaller) group of mapped IP addresses, on a first come, first served basis. Only the real host can initiate traffic. See [Dynamic NAT, on page 173](#).
- Dynamic Port Address Translation (PAT)—A group of real IP addresses are mapped to a single IP address using a unique source port of that IP address. See [Dynamic PAT, on page 178](#).

- Static NAT—A consistent mapping between a real and mapped IP address. Allows bidirectional traffic initiation. See [Static NAT, on page 182](#).
- Identity NAT—A real address is statically translated to itself, essentially bypassing NAT. You might want to configure NAT this way when you want to translate a large group of addresses, but then want to exempt a smaller subset of addresses. See [Identity NAT, on page 191](#).

NAT in Routed Mode

The following figure shows a typical NAT example in routed mode, with a private network on the inside.

Figure 5: NAT Example: Routed Mode



1. When the inside host at 10.1.2.27 sends a packet to a web server, the real source address of the packet, 10.1.2.27, is translated to a mapped address, 209.165.201.10.
2. When the server responds, it sends the response to the mapped address, 209.165.201.10, and the FTD device receives the packet because the FTD device performs proxy ARP to claim the packet.
3. The FTD device then changes the translation of the mapped address, 209.165.201.10, back to the real address, 10.1.2.27, before sending it to the host.

Auto NAT and Manual NAT

You can implement address translation in two ways: *auto NAT* and *manual NAT*.

We recommend using auto NAT unless you need the extra features that manual NAT provides. It is easier to configure auto NAT, and it might be more reliable for applications such as Voice over IP (VoIP). (For VoIP, you might see a failure in the translation of indirect addresses that do not belong to either of the objects used in the rule.)

Auto NAT

All NAT rules that are configured as a parameter of a network object are considered to be *auto NAT* rules. This is a quick and easy way to configure NAT for a network object. You cannot create these rules for a group object, however.

Although these rules are configured as part of the object itself, you cannot see the NAT configuration in the object definition through the object manager.

When a packet enters an interface, both the source and destination IP addresses are checked against the auto NAT rules. The source and destination address in the packet can be translated by separate rules if separate matches are made. These rules are not tied to each other; different combinations of rules can be used depending on the traffic.

Because the rules are never paired, you cannot specify that sourceA/destinationA should have a different translation than sourceA/destinationB. Use manual NAT for that kind of functionality, where you can identify the source and destination address in a single rule.

Manual NAT

Manual NAT lets you identify both the source and destination address in a single rule. Specifying both the source and destination addresses lets you specify that sourceA/destinationA can have a different translation than sourceA/destinationB.



Note For static NAT, the rule is bidirectional, so be aware that “source” and “destination” are used in commands and descriptions throughout this guide even though a given connection might originate at the “destination” address. For example, if you configure static NAT with port address translation, and specify the source address as a Telnet server, and you want all traffic going to that Telnet server to have the port translated from 2323 to 23, then you must specify the *source* ports to be translated (real: 23, mapped: 2323). You specify the source ports because you specified the Telnet server address as the source address.

The destination address is optional. If you specify the destination address, you can either map it to itself (identity NAT), or you can map it to a different address. The destination mapping is always a static mapping.

Comparing Auto NAT and Manual NAT

The main differences between these two NAT types are:

- How you define the real address.
 - Auto NAT—The NAT rule becomes a parameter for a network object. The network object IP address serves as the original (real) address.
 - Manual NAT—You identify a network object or network object group for both the real and mapped addresses. In this case, NAT is not a parameter of the network object; the network object or group is a parameter of the NAT configuration. The ability to use a network object *group* for the real address means that manual NAT is more scalable.
- How source and destination NAT is implemented.
 - Auto NAT— Each rule can apply to either the source or destination of a packet. So two rules might be used, one for the source IP address, and one for the destination IP address. These two rules cannot be tied together to enforce a specific translation for a source/destination combination.

- Manual NAT—A single rule translates both the source and destination. A packet matches one rule only, and further rules are not checked. Even if you do not configure the optional destination address, a matching packet still matches one manual NAT rule only. The source and destination are tied together, so you can enforce different translations depending on the source/destination combination. For example, sourceA/destinationA can have a different translation than sourceA/destinationB.
- Order of NAT Rules.
 - Auto NAT—Automatically ordered in the NAT table.
 - Manual NAT—Manually ordered in the NAT table (before or after auto NAT rules).

NAT Rule Order

Auto NAT and manual NAT rules are stored in a single table that is divided into three sections. Section 1 rules are applied first, then section 2, and finally section 3, until a match is found. For example, if a match is found in section 1, sections 2 and 3 are not evaluated. The following table shows the order of rules within each section.

Table 3: NAT Rule Table

Table Section	Rule Type	Order of Rules within the Section
Section 1	Manual NAT	<p>Applied on a first match basis, in the order they appear in the configuration. Because the first match is applied, you must ensure that specific rules come before more general rules, or the specific rules might not be applied as desired. By default, manual NAT rules are added to section 1.</p> <p>By "specific rules first," we mean:</p> <ul style="list-style-type: none"> • Static rules should come before dynamic rules. • Rules that include destination translation should come before rules with source translation only. <p>If you cannot eliminate overlapping rules, where more than one rule might apply based on the source or destination address, be especially careful to follow these recommendations.</p>

Table Section	Rule Type	Order of Rules within the Section
Section 2	Auto NAT	<p>If a match in section 1 is not found, section 2 rules are applied in the following order:</p> <ol style="list-style-type: none"> 1. Static rules. 2. Dynamic rules. <p>Within each rule type, the following ordering guidelines are used:</p> <ol style="list-style-type: none"> 1. Quantity of real IP addresses—From smallest to largest. For example, an object with one address will be assessed before an object with 10 addresses. 2. For quantities that are the same, then the IP address number is used, from lowest to highest. For example, 10.1.1.0 is assessed before 11.1.1.0. 3. If the same IP address is used, then the name of the network object is used, in alphabetical order. For example, abracadabra is assessed before catwoman.
Section 3	Manual NAT	<p>If a match is still not found, section 3 rules are applied on a first match basis, in the order they appear in the configuration. This section should contain your most general rules. You must also ensure that any specific rules in this section come before general rules that would otherwise apply.</p>

For section 2 rules, for example, you have the following IP addresses defined within network objects:

- 192.168.1.0/24 (static)
- 192.168.1.0/24 (dynamic)
- 10.1.1.0/24 (static)
- 192.168.1.1/32 (static)
- 172.16.1.0/24 (dynamic) (object def)
- 172.16.1.0/24 (dynamic) (object abc)

The resultant ordering would be:


- 192.168.1.1/32 (static)
- 10.1.1.0/24 (static)
- 192.168.1.0/24 (static)
- 172.16.1.0/24 (dynamic) (object abc)
- 172.16.1.0/24 (dynamic) (object def)
- 192.168.1.0/24 (dynamic)

NAT Interfaces

Except for bridge group member interfaces, you can configure a NAT rule to apply to any interface (in other words, all interfaces), or you can identify specific real and mapped interfaces. You can also specify any interface for the real address, and a specific interface for the mapped address, or vice versa.

For example, you might want to specify any interface for the real address and specify the outside interface for the mapped address if you use the same private addresses on multiple interfaces, and you want to translate them all to the same global pool when accessing the outside.

Figure 6: Specifying Any Interface



However, the concept of “any” interface does not apply to bridge group member interfaces. When you specify “any” interface, all bridge group member interfaces are excluded. Thus, to apply NAT to bridge group members, you must specify the member interface. This could result in many similar rules where only one interface is different. You cannot configure NAT for the Bridge Virtual Interface (BVI) itself, you can configure NAT for member interfaces only.

Configuring Routing for NAT

The FTD device needs to be the destination for any packets sent to the translated (mapped) address.

When sending packets, the device uses the destination interface if you specify one, or a routing table lookup if you do not, to determine the egress interface. For identity NAT, you have the option to use a route lookup even if you specify a destination interface.

The type of routing configuration needed depends on the type of mapped address, as explained in the following topics.

Addresses on the Same Network as the Mapped Interface

If you use addresses on the same network as the destination (mapped) interface, the FTD device uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the FTD device does not have to be the gateway for any additional networks. This solution is ideal if the outside network contains an adequate number of free addresses, a consideration if you are using a 1:1 translation like dynamic NAT or static NAT. Dynamic PAT greatly extends the number of translations you can use with a small number of addresses, so even if the available addresses on the outside network is small, this method can be used. For PAT, you can even use the IP address of the mapped interface.

Addresses on a Unique Network

If you need more addresses than are available on the destination (mapped) interface network, you can identify addresses on a different subnet. The upstream router needs a static route for the mapped addresses that points to the FTD device.

The Same Address as the Real Address (Identity NAT)

The default behavior for identity NAT has proxy ARP enabled, matching other static NAT rules. You can disable proxy ARP if desired. You can also disable proxy ARP for regular static NAT if desired, in which case you need to be sure to have proper routes on the upstream router.

Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues. For example, if you configure a broad identity NAT rule for “any” IP address, then leaving proxy ARP enabled can cause problems for hosts on the network directly connected to the mapped interface. In this case, when a host on the mapped network wants to communicate with another host on the same network, then the address in the ARP request matches the NAT rule (which matches “any” address). The FTD device will then proxy ARP for the address, even though the packet is not actually destined for the FTD device. (Note that this problem occurs even if you have a manual NAT rule; although the NAT rule must match both the source and destination addresses, the proxy ARP decision is made only on the “source” address). If the FTD device ARP response is received before the actual host ARP response, then traffic will be mistakenly sent to the FTD device.

Guidelines for NAT

The following topics provide detailed guidelines for implementing NAT.

Interface Guidelines

NAT is supported for standard routed physical or subinterfaces.

However, configuring NAT on bridge group member interfaces (interfaces that are part of a Bridge Virtual Interface, or BVI) has the following restrictions:

- When configuring NAT for the members of a bridge group, you specify the member interface. You cannot configure NAT for the bridge group interface (BVI) itself.
- When doing NAT between bridge group member interfaces, you must specify the source and destination interfaces. You cannot specify “any” as the interface.
- You cannot configure interface PAT when the destination interface is a bridge group member interface, because there is no IP address attached to the interface.
- You cannot translate between IPv4 and IPv6 networks (NAT64/46) when the source and destination interfaces are members of the same bridge group. Static NAT/PAT 44/66, dynamic NAT44/66, and dynamic PAT44 are the only allowed methods; dynamic PAT66 is not supported.

IPv6 NAT Guidelines

NAT supports IPv6 with the following guidelines and restrictions.

- For standard routed mode interfaces, you can also translate between IPv4 and IPv6.

- You cannot translate between IPv4 and IPv6 for interfaces that are members of the same bridge group. You can translate between two IPv6 or two IPv4 networks only. This restriction does not apply between a bridge group member and a standard routed interface.
- You cannot use dynamic PAT for IPv6 (NAT66) when translating between interfaces in the same bridge group. This restriction does not apply between a bridge group member and a standard routed interface.
- For static NAT, you can specify an IPv6 subnet up to /64. Larger subnets are not supported.
- When using FTP with NAT46, when an IPv4 FTP client connects to an IPv6 FTP server, the client must use either the extended passive mode (EPSV) or extended port mode (EPRT); PASV and PORT commands are not supported with IPv6.

IPv6 NAT Best Practices

You can use NAT to translate between IPv6 networks, and also to translate between IPv4 and IPv6 networks (routed mode only). We recommend the following best practices:

- NAT66 (IPv6-to-IPv6)—We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT. If you do not want to allow returning traffic, you can make the static NAT rule unidirectional (manual NAT only).
- NAT46 (IPv4-to-IPv6)—We recommend using static NAT. Because the IPv6 address space is so much larger than the IPv4 address space, you can easily accommodate a static translation. If you do not want to allow returning traffic, you can make the static NAT rule unidirectional (manual NAT only). When translating to an IPv6 subnet (/96 or lower), the resulting mapped address is by default an IPv4-embedded IPv6 address, where the 32-bits of the IPv4 address is embedded after the IPv6 prefix. For example, if the IPv6 prefix is a /96 prefix, then the IPv4 address is appended in the last 32-bits of the address. For example, if you map 192.168.1.0/24 to 201b::0/96, then 192.168.1.4 will be mapped to 201b::0.192.168.1.4 (shown with mixed notation). If the prefix is smaller, such as /64, then the IPv4 address is appended after the prefix, and a suffix of 0s is appended after the IPv4 address.
- NAT64 (IPv6-to-IPv4)—You may not have enough IPv4 addresses to accommodate the number of IPv6 addresses. We recommend using a dynamic PAT pool to provide a large number of IPv4 translations.

NAT Support for Inspected Protocols

Some application layer protocols that open secondary connections, or that embedded IP addresses in packets, are inspected to provide the following services:

- Pinhole creation—Some application protocols open secondary TCP or UDP connections either on standard or negotiated ports. Inspection opens pinholes for these secondary ports so that you do not need to create access control rules to allow them.
- NAT rewrite— Protocols such as FTP embed IP addresses and ports for the secondary connections in packet data as part of the protocol. If there is NAT translation involved for either of the endpoints, the inspection engines rewrite the packet data to reflect the NAT translation of the embedded addresses and ports. The secondary connections would not work without NAT rewrite.
- Protocol enforcement—Some inspections enforce some degree of conformance to the RFCs for the inspected protocol.

The following table lists the inspected protocols that apply NAT rewrite and their NAT limitations. Keep these limitations in mind when writing NAT rules that include these protocols. Inspected protocols not listed here do not apply NAT rewrite. These inspections include GTP, HTTP, IMAP, POP, SMTP, SSH, and SSL.



Note NAT rewrite is supported on the listed ports only. If you use these protocols on non-standard ports, do not use NAT on the connections.

Table 4: NAT Supported Application Inspection

Application	Inspected Protocol, Port	NAT Limitations	Pinholes Created
DCERPC	TCP/135	No NAT64.	Yes
DNS over UDP	UDP/53	No NAT support is available for name resolution through WINS.	No
ESMTP	TCP/25	No NAT64.	No
FTP	TCP/21	No limitations.	Yes
H.323 H.225 (Call signaling) H.323 RAS	TCP/1720 UDP/1718 For RAS, UDP/1718-1719	No NAT64.	Yes
ICMP ICMP Error	ICMP (ICMP traffic directed to a device interface is never inspected.)	No limitations.	No
IP Options	RSVP	No NAT64.	No
NetBIOS Name Server over IP	UDP/137, 138 (Source ports)	No NAT64.	No
RSH	TCP/514	No PAT. No NAT64.	Yes
RTSP	TCP/554 (No handling for HTTP cloaking.)	No NAT64.	Yes
SIP	TCP/5060 UDP/5060	No extended PAT. No NAT64 or NAT46.	Yes
Skinny (SCCP)	TCP/2000	No NAT64, NAT46, or NAT66.	Yes

Application	Inspected Protocol, Port	NAT Limitations	Pinholes Created
SQL*Net (versions 1, 2)	TCP/1521	No NAT64.	Yes
Sun RPC	TCP/111 UDP/111	No NAT64.	Yes
TFTP	UDP/69	No NAT64. Payload IP addresses are not translated.	Yes
XDMCP	UDP/177	No NAT64.	Yes

Additional Guidelines for NAT

- For interfaces that are members of a bridge group, you write NAT rules for the member interfaces. You cannot write NAT rules for the Bridge Virtual Interface (BVI) itself.
- (Auto NAT only.) You can only define a single NAT rule for a given object; if you want to configure multiple NAT rules for an object, you need to create multiple objects with different names that specify the same IP address.
- If a VPN is defined on an interface, inbound ESP traffic on the interface is not subject to the NAT rules. The system allows the ESP traffic for established VPN tunnels only, dropping traffic not associated with an existing tunnel. This restriction applies to ESP and UDP ports 500 and 4500.
- If you define a site-to-site VPN on a device that is behind a device that is applying dynamic PAT, so that UDP ports 500 and 4500 are not the ones actually used, you must initiate the connection from the device that is behind the PAT device. The responder cannot initiate the security association (SA) because it does not know the correct port numbers.
- If you change the NAT configuration, and you do not want to wait for existing translations to time out before the new NAT configuration is used, you can clear the translation table using the **clear xlate** command in the device CLI. However, clearing the translation table disconnects all current connections that use translations.

If you create a new NAT rule that should apply to an existing connection (such as a VPN tunnel), you need to use **clear conn** to end the connection. Then, the attempt to re-establish the connection should hit the NAT rule and the connection should be NAT'ed correctly.



Note If you remove a dynamic NAT or PAT rule, and then add a new rule with mapped addresses that overlap the addresses in the removed rule, then the new rule will not be used until all connections associated with the removed rule time out or are cleared using the **clear xlate** or **clear conn** commands. This safeguard ensures that the same address is not assigned to multiple hosts.



- You cannot use an object group with both IPv4 and IPv6 addresses; the object group must include only one type of address.

- (Manual NAT only.) When using **any** as the source address in a NAT rule, the definition of “any” traffic (IPv4 vs. IPv6) depends on the rule. Before the FTD device performs NAT on a packet, the packet must be IPv6-to-IPv6 or IPv4-to-IPv4; with this prerequisite, the FTD device can determine the value of **any** in a NAT rule. For example, if you configure a rule from “any” to an IPv6 server, and that server was mapped from an IPv4 address, then **any** means “any IPv6 traffic.” If you configure a rule from “any” to “any,” and you map the source to the interface IPv4 address, then **any** means “any IPv4 traffic” because the mapped interface address implies that the destination is also IPv4.
- You can use the same mapped object or group in multiple NAT rules.
- The mapped IP address pool cannot include:
 - The mapped interface IP address. If you specify “any” interface for the rule, then all interface IP addresses are disallowed. For interface PAT (routed mode only), specify the interface name instead of the interface address.
 - The failover interface IP address.
 - (Dynamic NAT.) The standby interface IP address when VPN is enabled.
- Avoid using overlapping addresses in static and dynamic NAT policies. For example, with overlapping addresses, a PPTP connection can fail to get established if the secondary connection for PPTP hits the static instead of dynamic xlate.
- You cannot use overlapping addresses in the source address of a NAT rule and a remote access VPN address pool.
- If you specify a destination interface in a rule, then that interface is used as the egress interface rather than looking up the route in the routing table. However, for identity NAT, you have the option to use a route lookup instead.
- NAT applies to through traffic only. Traffic generated by the system is not subject to NAT.
- Do not name a network object or group pat-pool, using any combination of upper- or lower-case letters.
- You cannot use NAT on the internal payload of Protocol Independent Multicast (PIM) registers.
- (Manual NAT) When writing NAT rules for a dual ISP interface setup (primary and backup interfaces using service level agreements in the routing configuration), do not specify destination criteria in the rule. Ensure the rule for the primary interface comes before the rule for the backup interface. This allows the device to choose the correct NAT destination interface based on the current routing state when the primary ISP is unavailable. If you specify destination objects, the NAT rule will always select the primary interface for the otherwise duplicate rules.
- If you get the ASP drop reason nat-no-xlate-to-pat-pool for traffic that should not match the NAT rules defined for the interface, configure identity NAT rules for the affected traffic so the traffic can pass untranslated.
- If you configure NAT for GRE tunnel endpoints, you must disable keepalives on the endpoints or the tunnel cannot be established. The endpoints send keepalives to the original addresses.

Configure NAT

Network address translation can be very complex. We recommend that you keep your rules as simple as possible to avoid translation problems and difficult troubleshooting situations. Careful planning before you implement NAT is critical. The following procedure provides the basic approach.

Procedure

- Step 1** Select **Policies > NAT**.
- Step 2** Decide what kinds of rules you need.
- You can create dynamic NAT, dynamic PAT, static NAT, and identity NAT rules. For an overview, see [NAT Types, on page 162](#).
- Step 3** Decide which rules should be implemented as manual or auto NAT.
- For a comparison of these two implementation options, see [Auto NAT and Manual NAT, on page 163](#).
- Step 4** Create the rules as explained in the following sections.
- [Dynamic NAT, on page 173](#)
 - [Dynamic PAT, on page 178](#)
 - [Static NAT, on page 182](#)
 - [Identity NAT, on page 191](#)
- Step 5** Manage the NAT policy and rules.
- You can do the following to manage the policy and its rules.
- To edit a rule, click the edit icon () for the rule.
 - To delete a rule, click the delete icon () for the rule.
-

Dynamic NAT

The following topics explain dynamic NAT and how to configure it.

About Dynamic NAT


Dynamic NAT translates a group of real addresses to a pool of mapped addresses that are routable on the destination network. The mapped pool typically includes fewer addresses than the real group. When a host you want to translate accesses the destination network, NAT assigns the host an IP address from the mapped pool. The translation is created only when the real host initiates the connection. The translation is in place only for the duration of the connection, and a given user does not keep the same IP address after the translation times out. Users on the destination network, therefore, cannot initiate a reliable connection to a host that uses dynamic NAT, even if the connection is allowed by an access rule.



Note For the duration of the translation, a remote host can initiate a connection to the translated host if an access rule allows it. Because the address is unpredictable, a connection to the host is unlikely. Nevertheless, in this case you can rely on the security of the access rule.


The following figure shows a typical dynamic NAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back.

Figure 7: Dynamic NAT



The following figure shows a remote host attempting to initiate a connection to a mapped address. This address is not currently in the translation table; therefore, the packet is dropped.

Figure 8: Remote Host Attempts to Initiate a Connection to a Mapped Address



Dynamic NAT Disadvantages and Advantages

Dynamic NAT has these disadvantages:

- If the mapped pool has fewer addresses than the real group, you could run out of addresses if the amount of traffic is more than expected.

Use PAT or a PAT fall-back method if this event occurs often because PAT provides over 64,000 translations using ports of a single address.

- You have to use a large number of routable addresses in the mapped pool, and routable addresses may not be available in large quantities.

The advantage of dynamic NAT is that some protocols cannot use PAT. PAT does not work with the following:

- IP protocols that do not have a port to overload, such as GRE version 0.
- Some multimedia applications that have a data stream on one port, the control path on another port, and are not open standard.

Configure Dynamic Auto NAT

Use dynamic auto NAT rules to translate addresses to different IP addresses that are routable on the destination network.

Before you begin

Select **Objects** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Address**—This must be a network object (not a group), and it can be a host or subnet.
- **Translated Address**—This can be a network object or group, but it cannot include a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

Procedure

-
- Step 1** Select **Policies > NAT**.
- Step 2** Do one of the following:
- To create a new rule, click the + button.
 - To edit an existing rule, click the edit icon (✎) for the rule.
- (To delete a rule you no longer need, click the trash can icon for the rule.)
- Step 3** Configure the basic rule options:
- **Title**—Enter a name for the rule.
 - **Create Rule For**—Select **Auto NAT**.
 - **Type**—Select **Dynamic**.
- Step 4** Configure the following packet translation options:
- **Source Interface, Destination Interface**—(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.
 - **Original Address**—The network object that contains the addresses you are translating.
 - **Translated Address**—The network object or group that contains the mapped addresses.

- Step 5** (Optional.) Click the **Advanced Options** link and select the desired options:
- **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see [Rewriting DNS Queries and Responses Using NAT, on page 232](#).
 - **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group.
- Step 6** Click **OK**.
-

Configure Dynamic Manual NAT

Use dynamic manual NAT rules when auto NAT does not meet your needs. For example, if you want to do different translations based on the destination. Dynamic NAT translates addresses to different IP addresses that are routable on the destination network.

Before you begin


Select **Objects** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source Address**—This can be a network object or group, and it can contain a host or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source Address**—This can be a network object or group, but it cannot include a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

You can also create network objects for the **Original Destination Address** and **Translated Destination Address** if you are configuring a static translation for those addresses in the rule.

For dynamic NAT, you can also perform port translation on the destination. In the Object Manager, ensure that there are port objects you can use for the **Original Destination Port** and **Translated Destination Port**. If you specify the source port, it will be ignored.

Procedure


- Step 1** Select **Policies > NAT**.
- Step 2** Do one of the following:
- To create a new rule, click the + button.
 - To edit an existing rule, click the edit icon () for the rule.
- (To delete a rule you no longer need, click the trash can icon for the rule.)

- Step 3** Configure the basic rule options:
- **Title**—Enter a name for the rule.
 - **Create Rule For**—Select **Manual NAT**.
 - **Rule Placement**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule you select.
 - **Type**—Select **Dynamic**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.

- Step 4** Configure the following interface options:
- **Source Interface, Destination Interface**—(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

- Step 5** Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet.



- **Original Source Address**—The network object or group that contains the addresses you are translating.
- **Original Destination Address**—(Optional.) The network object that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

You can select **Interface** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

- Step 6** Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.
- **Translated Source Address**—The network object or group that contains the mapped addresses.
 - **Translated Destination Address**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination Address**, you can set up identity NAT (that is, no translation) by selecting the same object.
- Step 7** (Optional.) Identify the destination service ports for service translation: **Original Destination Port, Translated Destination Port**.

Dynamic NAT does not support port translation, so leave the **Original Source Port** and **Translated Source Port** fields empty. However, because the destination translation is always static, you can perform port translation for the destination port.

NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

Step 8 (Optional.) Click the **Advanced Options** link and select the desired options:

- **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see [Rewriting DNS Queries and Responses Using NAT, on page 232](#).
- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group.

Step 9 Click **OK**.

Dynamic PAT

The following topics describe dynamic PAT.


About Dynamic PAT

Dynamic PAT translates multiple real addresses to a single mapped IP address by translating the real address and source port to the mapped address and a unique port.

Each connection requires a separate translation session because the source port differs for each connection. For example, 10.1.1.1:1025 requires a separate translation from 10.1.1.1:1026.

The following figure shows a typical dynamic PAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back. The mapped address is the same for each translation, but the port is dynamically assigned.

Figure 9: Dynamic PAT



For the duration of the translation, a remote host on the destination network can initiate a connection to the translated host if an access rule allows it. Because the port address (both real and mapped) is unpredictable, a connection to the host is unlikely. Nevertheless, in this case you can rely on the security of the access rule.

After the connection expires, the port translation also expires.



Note We recommend that you use different PAT pools for each interface. If you use the same pool for multiple interfaces, especially if you use it for "any" interface, the pool can be quickly exhausted, with no ports available for new translations.

Dynamic PAT Disadvantages and Advantages

Dynamic PAT lets you use a single mapped address, thus conserving routable addresses. You can even use the FTD device interface IP address as the PAT address. However, you cannot use interface PAT for the IPv6 addresses on the interface.

You cannot use dynamic PAT for IPv6 (NAT66) when translating between interfaces in the same bridge group. This restriction does not apply between a bridge group member and a standard routed interface.

Dynamic PAT does not work with some multimedia applications that have a data stream that is different from the control path. For more information, see [NAT Support for Inspected Protocols, on page 169](#).

Dynamic PAT might also create a large number of connections appearing to come from a single IP address, and servers might interpret the traffic as a DoS attack.

Configure Dynamic Auto PAT

Use dynamic auto PAT rules to translate addresses to unique IP address/port combinations, rather than to multiple IP addresses only. You can translate to a single address, either the destination interface's address or another address.

Before you begin

Select **Objects** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Address**—This must be a network object (not a group), and it can be a host or subnet.
- **Translated Address**—You have the following options to specify the PAT address:
 - **Destination Interface**—To use the destination interface IPv4 address, you do not need a network object. You cannot use interface PAT for IPv6.
 - **Single PAT address**—Create a network object containing a single host.

Procedure

Step 1 Select **Policies > NAT**.

Step 2 Do one of the following:

- To create a new rule, click the + button.

- To edit an existing rule, click the edit icon (✎) for the rule.

(To delete a rule you no longer need, click the trash can icon for the rule.)

Step 3 Configure the basic rule options:

- **Title**—Enter a name for the rule.
- **Create Rule For**—Select **Auto NAT**.
- **Type**—Select **Dynamic**.

Step 4 Configure the following packet translation options:

- **Source Interface, Destination Interface**—(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.
- **Original Address**—The network object that contains the addresses you are translating.
- **Translated Address**—One of the following:
 - (Interface PAT.) To use the IPv4 address of the destination interface, select **Interface**. You must also select a specific destination interface, which cannot be a bridge group member interface. You cannot use interface PAT for IPv6.
 - To use a single address other than the destination interface address, select the host network object you created for this purpose.

Step 5 (Optional.) Click the **Advanced Options** link and select the desired options:

- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. You cannot select this option if you already configured interface PAT as the translated address. You also cannot use this option with IPv6 networks.

Step 6 Click **OK**.

Configure Dynamic Manual PAT

Use dynamic manual PAT rules when auto PAT does not meet your needs. For example, if you want to do different translations based on the destination. Dynamic PAT translates addresses to unique IP address/port combinations, rather than to multiple IP addresses only. You can translate to a single address, either the destination interface's address or another address.

Before you begin

Select **Objects** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source Address**—This can be a network object or group, and it can contain a host or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source Address**—You have the following options to specify the PAT address:

- **Destination Interface**—To use the destination interface IPv4 address, you do not need a network object. You cannot use interface PAT for IPv6.
- **Single PAT address**—Create a network object containing a single host.


You can also create network objects for the **Original Destination Address** and **Translated Destination Address** if you are configuring a static translation for those addresses in the rule.

For dynamic PAT, you can also perform port translation on the destination. In the Object Manager, ensure that there are port objects you can use for the **Original Destination Port** and **Translated Destination Port**. If you specify the source port, it will be ignored.

Procedure

- Step 1** Select **Policies > NAT**.
- Step 2** Do one of the following:
- To create a new rule, click the + button.
 - To edit an existing rule, click the edit icon (✎) for the rule.
- (To delete a rule you no longer need, click the trash can icon for the rule.)
- Step 3** Configure the basic rule options:
- **Title**—Enter a name for the rule.
 - **Create Rule For**—Select **Manual NAT**.
 - **Rule Placement**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule you select.
 - **Type**—Select **Dynamic**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- Step 4** Configure the following interface options:
- **Source Interface, Destination Interface**—(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.
- Step 5** Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet.



- **Original Source Address**—The network object or group that contains the addresses you are translating.
- **Original Destination Address**—(Optional.) The network object that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

You can select **Interface** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source Address**—One of the following:
 - (Interface PAT.) To use the IPv4 address of the destination interface, select **Interface**. You must also select a specific destination interface, which cannot be a bridge group member interface. You cannot use interface PAT for IPv6.
 - To use a single address other than the destination interface address, select the host network object you created for this purpose.
- **Translated Destination Address**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the destination service ports for service translation: **Original Destination Port**, **Translated Destination Port**.

Dynamic NAT does not support port translation, so leave the **Original Source Port** and **Translated Source Port** fields empty. However, because the destination translation is always static, you can perform port translation for the destination port.

NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

Step 8 (Optional.) Click the **Advanced Options** link and select the desired options:

- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. You cannot select this option if you already configured interface PAT as the translated address. You also cannot use this option with IPv6 networks.

Step 9 Click **OK**.

Static NAT


The following topics explain static NAT and how to implement it.

About Static NAT

Static NAT creates a fixed translation of a real address to a mapped address. Because the mapped address is the same for each consecutive connection, static NAT allows bidirectional connection initiation, both to and from the host (if an access rule exists that allows it). With dynamic NAT and PAT, on the other hand, each host uses a different address or port for each subsequent translation, so bidirectional initiation is not supported.

The following figure shows a typical static NAT scenario. The translation is always active so both real and remote hosts can initiate connections.

Figure 10: Static NAT




Static NAT with Port Translation

Static NAT with port translation lets you specify a real and mapped protocol and port.

When you specify the port with static NAT, you can choose to map the port and/or the IP address to the same value or to a different value.

The following figure shows a typical static NAT with port translation scenario showing both a port that is mapped to itself and a port that is mapped to a different value; the IP address is mapped to a different value in both cases. The translation is always active so both translated and remote hosts can initiate connections.

Figure 11: Typical Static NAT with Port Translation Scenario



Static NAT-with-port-translation rules limit access to the destination IP address for the specified port only. If you try to access the destination IP address on a different port not covered by a NAT rule, then the connection is blocked. In addition, for manual NAT, traffic that does not match the source IP address of the NAT rule will be dropped if it matches the destination IP address, regardless of the destination port. Therefore, you must add additional rules for all other traffic allowed to the destination IP address. For example, you can configure a static NAT rule for the IP address, without port specification, and place it after the port translation rule.



Note For applications that require application inspection for secondary channels (for example, FTP and VoIP), NAT automatically translates the secondary ports.

Following are some other uses of static NAT with port translation.

Static NAT with Identity Port Translation

You can simplify external access to internal resources. For example, if you have three separate servers that provide services on different ports (such as FTP, HTTP, and SMTP), you can give external users a single IP address to access those services. You can then configure static NAT with identity port translation to map the single external IP address to the correct IP addresses of the real servers based on the port they are trying to access. You do not need to change the port, because the servers are using the standard ones (21, 80, and 25 respectively).

Static NAT with Port Translation for Non-Standard Ports

You can also use static NAT with port translation to translate a well-known port to a non-standard port or vice versa. For example, if inside web servers use port 8080, you can allow outside users to connect to port 80, and then undo translation to the original port 8080. Similarly, to provide extra security, you can tell web users to connect to non-standard port 6785, and then undo translation to port 80.

Static Interface NAT with Port Translation


You can configure static NAT to map a real address to an interface address/port combination. For example, if you want to redirect Telnet access for the device's outside interface to an inside host, then you can map the inside host IP address/port 23 to the outside interface address/port 23.

One-to-Many Static NAT

Typically, you configure static NAT with a one-to-one mapping. However, in some cases, you might want to configure a single real address to several mapped addresses (one-to-many). When you configure one-to-many static NAT, when the real host initiates traffic, it always uses the first mapped address. However, for traffic initiated to the host, you can initiate traffic to any of the mapped addresses, and they will be untranslated to the single real address.


The following figure shows a typical one-to-many static NAT scenario. Because initiation by the real host always uses the first mapped address, the translation of real host IP/first mapped IP is technically the only bidirectional translation.

Figure 12: One-to-Many Static NAT



For example, you have a load balancer at 10.1.2.27. Depending on the URL requested, it redirects traffic to the correct web server.

Figure 13: One-to-Many Static NAT Example




Other Mapping Scenarios (Not Recommended)

NAT has the flexibility to allow any kind of static mapping scenario: one-to-one, one-to-many, but also few-to-many, many-to-few, and many-to-one mappings. We recommend using only one-to-one or one-to-many mappings. These other mapping options might result in unintended consequences.

Functionally, few-to-many is the same as one-to-many; but because the configuration is more complicated and the actual mappings may not be obvious at a glance, we recommend creating a one-to-many configuration for each real address that requires it. For example, for a few-to-many scenario, the few real addresses are mapped to the many mapped addresses in order (A to 1, B to 2, C to 3). When all real addresses are mapped, the next mapped address is mapped to the first real address, and so on until all mapped addresses are mapped (A to 4, B to 5, C to 6). This results in multiple mapped addresses for each real address. Just like a one-to-many configuration, only the first mappings are bidirectional; subsequent mappings allow traffic to be initiated *to* the real host, but all traffic *from* the real host uses only the first mapped address for the source.

The following figure shows a typical few-to-many static NAT scenario.

Figure 14: Few-to-Many Static NAT



For a many-to-few or many-to-one configuration, where you have more real addresses than mapped addresses, you run out of mapped addresses before you run out of real addresses. Only the mappings between the lowest real IP addresses and the mapped pool result in bidirectional initiation. The remaining higher real addresses can initiate traffic, but traffic cannot be initiated to them (returning traffic for a connection is directed to the correct real address because of the unique 5-tuple (source IP, destination IP, source port, destination port, protocol) for the connection).



Note Many-to-few or many-to-one NAT is not PAT. If two real hosts use the same source port number and go to the same outside server and the same TCP destination port, and both hosts are translated to the same IP address, then both connections will be reset because of an address conflict (the 5-tuple is not unique).

The following figure shows a typical many-to-few static NAT scenario.

Figure 15: Many-to-Few Static NAT



Instead of using a static rule this way, we suggest that you create a one-to-one rule for the traffic that needs bidirectional initiation, and then create a dynamic rule for the rest of your addresses.

Configure Static Auto NAT

Use static auto NAT rules to translate addresses to different IP addresses that are routable on the destination network. You can also do port translation with the static NAT rule.

Before you begin

Select **Objects** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Address**—This must be a network object (not a group), and it can be a host or subnet.

- **Translated Address**—You have the following options to specify the translated address:
 - **Destination Interface**—To use the destination interface IPv4 address, you do not need a network object. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number. You cannot use interface PAT for IPv6.
 - **Address**—Create a network object or group containing hosts or subnets. A group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.

Procedure

- Step 1** Select **Policies > NAT**.
- Step 2** Do one of the following:
- To create a new rule, click the + button.
 - To edit an existing rule, click the edit icon (✎) for the rule.
- (To delete a rule you no longer need, click the trash can icon for the rule.)
- Step 3** Configure the basic rule options:
- **Title**—Enter a name for the rule.
 - **Create Rule For**—Select **Auto NAT**.
 - **Type**—Select **Static**.
- Step 4** Configure the following packet translation options:
- **Source Interface, Destination Interface**—(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.
 - **Original Address**—The network object that contains the addresses you are translating.
 - **Translated Address**—One of the following:
 - To use a set group of addresses, select the network object or group that contains the mapped addresses. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.
 - (Static interface NAT with port translation.) To use the address of the destination interface, select **Interface**. You must also select a specific destination interface, which cannot be a bridge group member interface. You cannot use interface PAT for IPv6. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
 - (Optional.) **Original Port, Translated Port**—If you need to translate a TCP or UDP port, select the port objects that define the original and translated ports. The objects must be for the same protocol. Click the **Create New Object** link if the objects do not already exist. For example, you can translate TCP/80 to TCP/8080 if necessary.
- Step 5** (Optional.) Click the **Advanced Options** link and select the desired options:

- **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see [Rewriting DNS Queries and Responses Using NAT, on page 232](#). This option is not available if you are doing port translation.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.

Step 6 Click **OK**.

Configure Static Manual NAT

Use static manual NAT rules when auto NAT does not meet your needs. For example, if you want to do different translations based on the destination. Static NAT translates addresses to different IP addresses that are routable on the destination network. You can also do port translation with the static NAT rule.

Before you begin

Select **Objects** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:


- **Original Source Address**—This can be a network object or group, and it can contain a host or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source Address**—You have the following options to specify the translated address:
 - **Destination Interface**—To use the destination interface IPv4 address, you do not need a network object. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number. You cannot use interface PAT for IPv6.
 - **Address**—Create a network object or group containing hosts or subnets. A group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.

You can also create network objects for the **Original Destination Address** and **Translated Destination Address** if you are configuring a static translation for those addresses in the rule. If you want to configure destination static interface NAT with port translation only, you can skip adding an object for the destination mapped addresses and specify the interface in the rule.

You can also perform port translation on the source, destination, or both. In the Object Manager, ensure that there are port objects you can use for the original and translated ports.

Procedure

- Step 1** Select **Policies > NAT**.
- Step 2** Do one of the following:
- To create a new rule, click the + button.
 - To edit an existing rule, click the edit icon (✎) for the rule.
- (To delete a rule you no longer need, click the trash can icon for the rule.)
- Step 3** Configure the basic rule options:
- Title**—Enter a name for the rule.
 - Create Rule For**—Select **Manual NAT**.
 - Rule Placement**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule you select.
 - Type**—Select **Static**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- Step 4** Configure the following interface options:
- Source Interface, Destination Interface**—(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.
- Step 5** Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.
- See the following figure for an example of the original packet vs. the translated packet.



- Original Source Address**—The network object or group that contains the addresses you are translating.
- Original Destination Address**—(Optional.) The network object that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

You can select **Interface** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source Address**—One of the following:
 - To use a set group of addresses, select the network object or group that contains the mapped addresses. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.
 - (Static interface NAT with port translation.) To use the IPv4 address of the destination interface, select **Interface**. You must also select a specific destination interface, which cannot be a bridge group member interface. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number. You cannot use interface PAT for IPv6.
- **Translated Destination Address**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the source or destination service ports for service translation.

If you are configuring static NAT with port translation, you can translate ports for the source, destination, or both. For example, you can translate between TCP/80 and TCP/8080.

NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

- **Original Source Port, Translated Source Port**—Defines a port translation for the source address.
- **Original Destination Port, Translated Destination Port**—Defines a port translation for the destination address.

Step 8 (Optional.) Click the **Advanced Options** link and select the desired options:

- **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see [Rewriting DNS Queries and Responses Using NAT, on page 232](#). This option is not available if you are doing port translation.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.


Step 9 Click **OK**.

Identity NAT

You might have a NAT configuration in which you need to translate an IP address to itself. For example, if you create a broad rule that applies NAT to every network, but want to exclude one network from NAT, you can create a static NAT rule to translate an address to itself.

The following figure shows a typical identity NAT scenario.

Figure 16: Identity NAT



The following topics explain how to configure identity NAT.

Configure Identity Auto NAT

Use static identity auto NAT rules to prevent the translation of an address. That is, to translate the address to itself.

Before you begin

Select **Objects** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Address**—This must be a network object (not a group), and it can be a host or subnet.
- **Translated Address**—A network object or group with the exact same contents as the original source object. You can use the same object.

Procedure

-
- Step 1** Select **Policies > NAT**.
- Step 2** Do one of the following:
- To create a new rule, click the + button.
 - To edit an existing rule, click the edit icon (✎) for the rule.
- (To delete a rule you no longer need, click the trash can icon for the rule.)
- Step 3** Configure the basic rule options:
- **Title**—Enter a name for the rule.
 - **Create Rule For**—Select **Auto NAT**.
 - **Type**—Select **Static**.
- Step 4** Configure the following packet translation options:

- **Source Interface, Destination Interface**—(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.
- **Original Address**—The network object that contains the addresses you are translating.
- **Translated Address**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

Do not configure the **Original Port** and **Translated Port** options for identity NAT.

Step 5 (Optional.) Click the **Advanced Options** link and select the desired options:

- **Translate DNS replies that match this rule**—Do not configure this option for identity NAT.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.
- **Perform Route Lookup for Destination Interface**— If you select source and destination interfaces when selecting the same object for original and translated source address, you can select this option to have the system determine the destination interface based on the routing table rather than using the destination interface configured in the NAT rule.

Step 6 Click **OK**.

Configure Identity Manual NAT

Use static identity manual NAT rules when auto NAT does not meet your needs. For example, if you want to do different translations based on the destination. Use static identity NAT rules to prevent the translation of an address. That is, to translate the address to itself.

Before you begin

Select **Objects** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source Address**—This can be a network object or group, and it can contain a host or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source Address**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.


You can also create network objects for the **Original Destination Address** and **Translated Destination Address** if you are configuring a static translation for those addresses in the rule. If you want to configure destination static interface NAT with port translation only, you can skip adding an object for the destination mapped addresses and specify the interface in the rule.

You can also perform port translation on the source, destination, or both. In the Object Manager, ensure that there are port objects you can use for the original and translated ports. You can use the same object for identity NAT.

Procedure

- Step 1** Select **Policies > NAT**.
- Step 2** Do one of the following:
- To create a new rule, click the + button.
 - To edit an existing rule, click the edit icon (✎) for the rule.
- (To delete a rule you no longer need, click the trash can icon for the rule.)
- Step 3** Configure the basic rule options:
- Title**—Enter a name for the rule.
 - Create Rule For**—Select **Manual NAT**.
 - Rule Placement**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule you select.
 - Type**—Select **Static**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- Step 4** Configure the following interface options:
- Source Interface, Destination Interface**—(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.
- Step 5** Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet where you perform identity NAT on the inside host but translate the outside host.



- Original Source Address**—The network object or group that contains the addresses you are translating.
- Original Destination Address**—(Optional.) The network object that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

You can select **Interface** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source Address**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.
- **Translated Destination Address**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination Address**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the source or destination service ports for service translation.

If you are configuring static NAT with port translation, you can translate ports for the source, destination, or both. For example, you can translate between TCP/80 and TCP/8080.

NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

- **Original Source Port, Translated Source Port**—Defines a port translation for the source address.
- **Original Destination Port, Translated Destination Port**—Defines a port translation for the destination address.

Step 8 (Optional.) Click the **Advanced Options** link and select the desired options:

- **Translate DNS replies that match this rule**—Do not configure this option for identity NAT.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.
- **Perform route lookup for Destination interface**— If you select source and destination interfaces when selecting the same object for original and translated source address, you can select this option to have the system determine the destination interface based on the routing table rather than using the destination interface configured in the NAT rule.

Step 9 Click **OK**.

NAT Rule Properties for FTD

Use Network Address Translation (NAT) rules to translate IP addresses to other IP addresses. You would typically use NAT rules to convert private addresses to publically routable addresses. The translation can be from one address to another, or you can use Port Address Translation (PAT) to translate many addresses to one, using port numbers to distinguish among the source addresses.

NAT rules include the following basic properties. The properties are the same for auto NAT and manual NAT rules except where indicated.

Title

Enter a name for the rule. The name cannot include spaces.

Create Rule For

Whether the translation rule is **Auto NAT** or **Manual NAT**. Auto NAT is simpler than manual NAT, but manual NAT allows you to create separate translations for a source address based on the destination address.

Status

Whether you want the rule to be active or disabled.

Placement (Manual NAT only.)

Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule you select.

Type

Whether the translation rule is **Dynamic** or **Static**. Dynamic translation automatically chooses the mapped address from a pool of addresses, or an address/port combination when implementing PAT. Use static translation if you want to precisely define the mapped address/port.

The following topics describe the remaining NAT rules properties.

Packet Translation Properties for Auto NAT

Use the **Packet Translation** options to define the source addresses and the mapped translated addresses. The following properties apply to auto NAT only.

Source Interface, Destination Interface

(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Original Address (Always required.)

The network object that contains the source addresses you are translating. This must be a network object (not a group), and it can be a host or subnet.

Translated Address (Usually required.)

The mapped addresses, the ones to which you are translating. What you select here depends on the type of translation rule you are defining.

- **Dynamic NAT**—The network object or group that contains the mapped addresses. This can be a network object or group, but it cannot include a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.
- **Dynamic PAT**—One of the following:
 - (Interface PAT.) To use the IPv4 address of the destination interface, select **Interface**. You must also select a specific destination interface, which cannot be a bridge group member interface. You cannot use interface PAT for IPv6.
 - To use a single address other than the destination interface address, select the host network object you created for this purpose.
- **Static NAT**—One of the following:

- To use a set group of addresses, select the network object or group that contains the mapped addresses. The object or group can contain hosts or subnets. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.
- (Static interface NAT with port translation.) To use the address of the destination interface, select **Interface**. You must also select a specific destination interface, which cannot be a bridge group member interface. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number. You cannot use interface PAT for IPv6.
- **Identity NAT**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

Original Port, Translated Port (Static NAT only.)

If you need to translate a TCP or UDP port, select the port objects that define the original and translated ports. The objects must be for the same protocol. For example, you can translate TCP/80 to TCP/8080 if necessary.

Packet Translation Properties for Manual NAT

Use the **Packet Translation** options to define the source addresses and the mapped translated addresses. The following properties apply to manual NAT only. All are optional except as indicated.

Source Interface, Destination Interface

(Required for bridge group member interfaces.) The interfaces where this NAT rule applies. **Source** is the real interface, the one through which the traffic enters the device. **Destination** is the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Original Source Address (Always required.)

The network object or group that contains the addresses you are translating. This can be a network object or group, and it can contain a host or subnet. If you want to translate all original source traffic, you can specify **Any** in the rule.

Translated Source Address (Usually required.)

The mapped addresses, the ones to which you are translating. What you select here depends on the type of translation rule you are defining.

- **Dynamic NAT**—The network object or group that contains the mapped addresses. This can be a network object or group, but it cannot include a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.
- **Dynamic PAT**—One of the following:
 - (Interface PAT.) To use the address of the destination interface, select **Interface**. You must also select a specific destination interface, which cannot be a bridge group member interface. You cannot use interface PAT for IPv6.
 - To use a single address other than the destination interface address, select the host network object you created for this purpose.
- **Static NAT**—One of the following:

- To use a set group of addresses, select the network object or group that contains the mapped addresses. The object or group can contain hosts or subnets. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.
- (Static interface NAT with port translation.) To use the address of the destination interface, select **Interface**. You must also select a specific destination interface, which cannot be a bridge group member interface. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number. You cannot use interface PAT for IPv6.
- **Identity NAT**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

Original Destination Address

The network object that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

You can select **Interface** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Translated Destination Address

The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Original Source Port, Translated Source Port, Original Destination Port, Translated Destination Port

The port objects that define the source and destination services for the original and translated packets. You can translate the ports, or select the same object to make the rule sensitive to the service without translating the ports. Keep the following rules in mind when configuring services:

- (Dynamic NAT or PAT.) You cannot do translation on the **Original Source Port** and **Translated Source Port**. You can do translation on the destination port only.
- NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same object for both the real and mapped ports.

Advanced NAT Properties

When you configure NAT, you can configure properties that provide specialized services in the **Advanced** options. All of these properties are optional: configure them only if you need the service.

Translate DNS replies that match this rule

Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and

AAAA records. For more information, see [Rewriting DNS Queries and Responses Using NAT, on page 232](#). This option is not available if you are doing port translation in a static NAT rule.

Fallthrough to Interface PAT (Destination Interface) (Dynamic NAT only.)

Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. You cannot select this option if you already configured interface PAT as the translated address. You cannot use this option with IPv6 networks.

Do not proxy ARP on Destination Interface (Static NAT only.)

Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.

Perform Route Lookup for Destination Interface (Static Identity NAT only. Routed mode only.)

If you select source and destination interfaces when selecting the same object for original and translated source address, you can select this option to have the system determine the destination interface based on the routing table rather than using the destination interface configured in the NAT rule.

Translating IPv6 Networks

In cases where you need to pass traffic between IPv6-only and IPv4-only networks, you need to use NAT to convert between the address types. Even with two IPv6 networks, you might want to hide internal addresses from the outside network.

You can use the following translation types with IPv6 networks:

- NAT64, NAT46—Translates IPv6 packets into IPv4 and vice versa. You need to define two policies, one for the IPv6 to IPv4 translation, and one for the IPv4 to IPv6 translation. Although you can accomplish this with a single manual NAT rule, if the DNS server is on the external network, you probably need to rewrite the DNS response. Because you cannot enable DNS rewrite on a manual NAT rule when you specify a destination, creating two auto NAT rules is the better solution.



Note NAT46 supports static mappings only.

- NAT66—Translates IPv6 packets to a different IPv6 address. We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT.



Note NAT64 and NAT 46 are possible on standard routed interfaces only. NAT66 is possible on both routed and bridge group member interfaces.

NAT64/46: Translating IPv6 Addresses to IPv4


When traffic goes from an IPv6 network to an IPv4-only network, you need to convert the IPv6 address to IPv4, and return traffic from IPv4 to IPv6. You need to define two address pools, an IPv4 address pool to bind IPv6 addresses in the IPv4 network, and an IPv6 address pool to bind IPv4 addresses in the IPv6 network.

- The IPv4 address pool for the NAT64 rule is normally small and typically might not have enough addresses to map one-to-one with the IPv6 client addresses. Dynamic PAT might more easily meet the possible large number of IPv6 client addresses compared to dynamic or static NAT.
- The IPv6 address pool for the NAT46 rule can be equal to or larger than the number of IPv4 addresses to be mapped. This allows each IPv4 address to be mapped to a different IPv6 address. NAT46 supports static mappings only, so you cannot use dynamic PAT.

You need to define two policies, one for the source IPv6 network, and one for the destination IPv4 network. Although you can accomplish this with a single manual NAT rule, if the DNS server is on the external network, you probably need to rewrite the DNS response. Because you cannot enable DNS rewrite on a manual NAT rule when you specify a destination, creating two auto NAT rules is the better solution.

NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet

Following is a straight-forward example where you have an inside IPv6-only network, and you want to convert to IPv4 for traffic sent to the Internet. This example assumes you do not need DNS translation, so you can perform both the NAT64 and NAT46 translations in a single manual NAT rule.



In this example, you translate the inside IPv6 network to IPv4 using dynamic interface PAT with the IP address of the outside interface. Outside IPv4 traffic is statically translated to addresses on the 2001:db8::/96 network, allowing transmission on the inside network.

Procedure

- Step 1** Create a network object for the inside IPv6 network.
- Choose **Objects**.
 - Select **Network** from the table of contents and click +.
 - Define the inside IPv6 network.

Name the network object (for example, `inside_v6`), select **Network**, and enter the network address, `2001:db8::/96`.

Add Network Object

Name
inside_v6

Description

Type
 Network Host

Network
2001:DB8::/96

d) Click **OK**.

Step 2

Create the manual NAT rule to translate the IPv6 network to IPv4 and back again.

- a) Select **Policies > NAT**.
- b) Click the + button.
- c) Configure the following properties:
 - **Title** = PAT64Rule (or another name of your choosing).
 - **Create Rule For** = Manual NAT.
 - **Placement** = Before Auto NAT Rules
 - **Type** = Dynamic.
 - **Source Interface** = inside.
 - **Destination Interface** = outside.
 - **Original Packet Source Address** = inside_v6 network object.
 - **Translated Packet Source Address** = Interface. This option uses the IPv4 address of the destination interface as the PAT address.
 - **Original Packet Destination Address** = inside_v6 network object.
 - **Translated Packet Destination Address** = any-ipv4 network object.

Title	Create Rule for	Status
PAT64Rule	Manual NAT	<input checked="" type="checkbox"/>

Manual NAT rules allow the translation of the source as well as the destination address of a network packet. Destination and port translation are optional. You can place manual NAT rules either before or after Auto NAT rules and insert the rules at a specific location.

Placement	Type
Before Auto NAT Rules	Dynamic

Packet Translation

ORIGINAL PACKET

Source Interface: inside

Source Address: inside_v6

Source Port: Any

Destination Address: inside_v6

Destination Port: Any

TRANSLATED PACKET

Destination Interface: outside

Source Address: Interface

Source Port: Any

Destination Address: any-ipv4


Destination Port: Any

d) Click **OK**.

With this rule, any traffic from the 2001:db8::/96 subnet on the inside interface going to the outside interface gets a NAT64 PAT translation using the IPv4 address of the outside interface. Conversely, any IPv4 address on the outside network coming to the inside interface is translated to an address on the 2001:db8::/96 network using the embedded IPv4 address method.

NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet and DNS Translation

Following is a typical example where you have an inside IPv6-only network, but there are some IPv4-only services on the outside Internet that internal users need.



In this example, you translate the inside IPv6 network to IPv4 using dynamic interface PAT with the IP address of the outside interface. Outside IPv4 traffic is statically translated to addresses on the 2001:db8::/96 network, allowing transmission on the inside network. You enable DNS rewrite on the NAT46 rule, so that replies from the external DNS server can be converted from A (IPv4) to AAAA (IPv6) records, and the addresses converted from IPv4 to IPv6.

Following is a typical sequence for a web request where a client at 2001:DB8::100 on the internal IPv6 network tries to open www.example.com.

1. The client's computer sends a DNS request to the DNS server at 2001:DB8::D1A5:CA81. The NAT rules make the following translations to the source and destination in the DNS request:
 - 2001:DB8::100 to a unique port on 209.165.201.1 (The NAT64 interface PAT rule.)
 - 2001:DB8::D1A5:CA81 to 209.165.202.129 (The NAT46 rule. D1A5:CA81 is the IPv6 equivalent of 209.165.202.129.)
2. The DNS server responds with an A record indicating that www.example.com is at 209.165.200.225. The NAT46 rule, with DNS rewrite enabled, converts the A record to the IPv6-equivalent AAAA record, and translates 209.165.200.225 to 2001:db8:D1A5:C8E1 in the AAAA record. In addition, the source and destination addresses in the DNS response are untranslated:
 - 209.165.202.129 to 2001:DB8::D1A5:CA81
 - 209.165.201.1 to 2001:db8::100
3. The IPv6 client now has the IP address of the web server, and makes an HTTP request to www.example.com at 2001:db8:D1A5:C8E1. (D1A5:C8E1 is the IPv6 equivalent of 209.165.200.225.) The source and destination of the HTTP request are translated:
 - 2001:DB8::100 to a unique port on 209.156.101.54 (The NAT64 interface PAT rule.)
 - 2001:db8:D1A5:C8E1 to 209.165.200.225 (The NAT46 rule.)

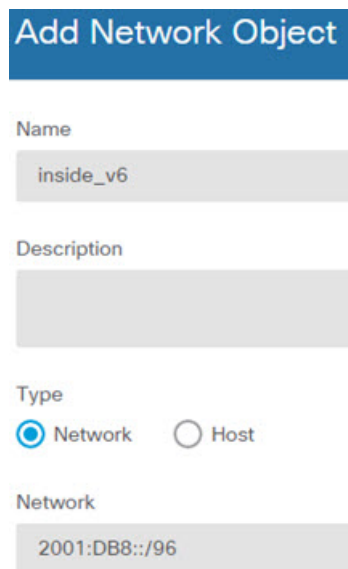
The following procedure explains how to configure this example.

Procedure

Step 1 Create the network objects that define the inside IPv6 and outside IPv4 networks.

- a) Choose **Objects**.
- b) Select **Network** from the table of contents and click +.
- c) Define the inside IPv6 network.

Name the network object (for example, `inside_v6`), select **Network**, and enter the network address, `2001:db8::/96`.

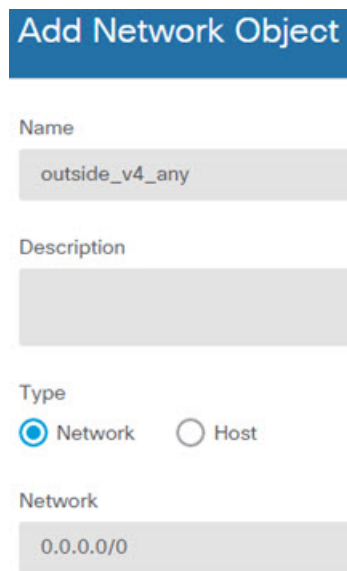


The screenshot shows a configuration form titled "Add Network Object". It contains the following fields and options:

- Name:** A text input field containing "inside_v6".
- Description:** An empty text input field.
- Type:** Two radio button options: "Network" (which is selected) and "Host".
- Network:** A text input field containing "2001:DB8::/96".

- d) Click **OK**.
- e) Click + and define the outside IPv4 network.

Name the network object (for example, `outside_v4_any`), select **Network**, and enter the network address `0.0.0.0/0`.



Add Network Object

Name
outside_v4_any

Description

Type
 Network Host

Network
0.0.0.0/0

Step 2 Configure the NAT64 dynamic PAT rule for the inside IPv6 network.

- a) Select **Policies > NAT**.
- b) Click the + button.
- c) Configure the following properties:
 - **Title** = PAT64Rule (or another name of your choosing).
 - **Create Rule For** = Auto NAT.
 - **Type** = Dynamic.
 - **Source Interface** = inside.
 - **Destination Interface** = outside.
 - **Original Address** = inside_v6 network object.
 - **Translated Address = Interface**. This option uses the IPv4 address of the destination interface as the PAT address.

Add NAT Rule

Title PAT64Rule **Create Rule for** Auto NAT **Status**

Auto NAT rules translate a specified host or network address regardless of its appearance as the source or destination address of a packet. These rules are automatically ordered and placed in the Auto NAT section.

Placement Automatically placed in Auto NAT rules **Type** Dynamic

Packet Translation | **Advanced Options**

ORIGINAL PACKET		TRANSLATED PACKET	
Source Interface	inside	Destination Interface	outside
Original Address	inside_v6	Translated Address	Interface
Original Port	Any	Translated Port	Any

d) Click **OK**.

With this rule, any traffic from the 2001:db8::/96 subnet on the inside interface going to the outside interface gets a NAT64 PAT translation using the IPv4 address of the outside interface.

Step 3 Configure the static NAT46 rule for the outside IPv4 network.

- Click the + button.
- Configure the following properties:
 - **Title** = NAT46Rule (or another name of your choosing).
 - **Create Rule For** = Auto NAT.
 - **Type** = Static.
 - **Source Interface** = outside.
 - **Destination Interface** = inside.
 - **Original Address** = outside_v4_any network object.
 - **Translated Address** = inside_v6 network object.
 - On the **Advanced Options** tab, select **Translate DNS replies that match this rule**.

Add NAT Rule ?

Title

Create Rule for

Auto NAT
▼

Status

Auto NAT rules translate a specified host or network address regardless of its appearance as the source or destination address of a packet. These rules are automatically ordered and placed in the Auto NAT section.

Placement

Automatically placed in Auto NAT rules

Type

Static
▼

Packet Translation

Advanced Options

ORIGINAL PACKET

Source Interface

outside
▼

Original Address

outside_v4_any
▼

Original Port

Any
▼

TRANSLATED PACKET

Destination Interface

inside

Translated Address

inside_v6
▼

Translated Port

Any

c) Click **OK**.

With this rule, any IPv4 address on the outside network coming to the inside interface is translated to an address on the 2001:db8::/96 network using the embedded IPv4 address method. In addition, DNS responses are converted from A (IPv4) to AAAA (IPv6) records, and the addresses converted from IPv4 to IPv6.


NAT66: Translating IPv6 Addresses to Different IPv6 Addresses

When going from an IPv6 network to another IPv6 network, you can translate the addresses to different IPv6 addresses on the outside network. We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT.

Because you are not translating between different address types, you need a single rule for NAT66 translations. You can easily model these rules using auto NAT. However, if you do not want to allow returning traffic, you can make the static NAT rule unidirectional using manual NAT only.

NAT66 Example, Static Translation between Networks

You can configure a static translation between IPv6 address pools using auto NAT. The following example explains how to convert inside addresses on the 2001:db8:122:2091::/96 network to outside addresses on the 2001:db8:122:2999::/96 network.



Note This example assumes that the inside interface is not a bridge group interface (BVI) but a standard routed interface. If the inside interface is a BVI, you need to duplicate the rules for each member interface.

Procedure

Step 1 Create the network objects that define the inside IPv6 and outside IPv6 NAT networks.

- a) Choose **Objects**.
- b) Select **Network** from the table of contents and click +.
- c) Define the inside IPv6 network.

Name the network object (for example, `inside_v6`), select **Network**, and enter the network address, `2001:db8:122:2091::/96`.

Add Network Object

Name

inside_v6

Description

Type



Network



Host

Network

2001:db8:122:2091::/96

- d) Click **OK**.
- e) Click + and define the outside IPv6 NAT network.

Name the network object (for example, outside_nat_v6), select **Network**, and enter the network address 2001:db8:122:2999::/96.

Add Network Object

Name

outside_nat_v6

Description

Type



Network



Host

Network

2001:db8:122:2999::/96

Step 2 Configure the static NAT rule for the inside IPv6 network.

- a) Select **Policies > NAT**.
- b) Click the + button.
- c) Configure the following properties:
 - **Title** = NAT66Rule (or another name of your choosing).
 - **Create Rule For** = Auto NAT.

- **Type** = Static.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Address** = inside_v6 network object.
- **Translated Address** = outside_nat_v6 network object.

Add NAT Rule ?

Title NAT66Rule	Create Rule for Auto NAT	Status <input checked="" type="checkbox"/>
---------------------------	------------------------------------	------------------------------------------------------

Auto NAT rules translate a specified host or network address regardless of its appearance as the source or destination address of a packet. These rules are automatically ordered and placed in the Auto NAT section.

Placement Automatically placed in Auto NAT rules	Type Static
------------------------------------------------------------	-----------------------

Packet Translation

Advanced Options

ORIGINAL PACKET		TRANSLATED PACKET	
Source Interface inside		Destination Interface outside	
Original Address inside_v6	Original Port Any	Translated Address outside_nat_v6	Translated Port Any


d) Click **OK**.

With this rule, any traffic from the 2001:db8:122:2091::/96 subnet on the inside interface going to the outside interface gets a static NAT66 translation to an address on the 2001:db8:122:2999::/96 network.

NAT66 Example, Simple IPv6 Interface PAT

A simple approach for implementing NAT66 is to dynamically assign internal addresses to different ports on the outside interface IPv6 address.

However, you cannot configure interface PAT using the IPv6 address of an interface using the FDM. Instead, use a single free address on the same network as a dynamic PAT pool.



Note This example assumes that the inside interface is not a bridge group interface (BVI) but a standard routed interface. If the inside interface is a BVI, you need to duplicate the rules for each member interface.

Procedure

Step 1 Create the network objects that define the inside IPv6 network and the IPv6 PAT address.

- a) Choose **Objects**.
- b) Select **Network** from the table of contents and click +.
- c) Define the inside IPv6 network.

Name the network object (for example, `inside_v6`), select **Network**, and enter the network address, `2001:db8:122:2091::/96`.

Add Network Object

Name
inside_v6

Description

Type
 Network Host

Network
2001:db8:122:2091::/96

- d) Click **OK**.
- e) Click + and define the outside IPv6 PAT address.

Name the network object (for example, ipv6_pat), select **Host**, and enter the host address 2001:db8:122:201b::2.

Add Network Object

Name
ipv6_pat

Description

Type
 Network Host

Host
2001:db8:122:201b::2

Step 2 Configure the dynamic PAT rule for the inside IPv6 network.

- a) Select **Policies > NAT**.
- b) Click the + button.
- c) Configure the following properties:
 - **Title** = PAT66Rule (or another name of your choosing).
 - **Create Rule For** = Auto NAT.

- **Type** = Dynamic.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Address** = inside_v6 network object.
- **Translated Address** = ipv6_pat network object.

Add NAT Rule ?

Title	Create Rule for	Status
<input type="text" value="PAT66Rule"/>	<input style="border: none; background-color: #f0f0f0; border-radius: 4px; padding: 2px 10px;" type="text" value="Auto NAT"/> ▼	<input checked="" type="checkbox"/>

Auto NAT rules translate a specified host or network address regardless of its appearance as the source or destination address of a packet. These rules are automatically ordered and placed in the Auto NAT section.

Placement	Type
<input type="text" value="Automatically placed in Auto NAT rules"/>	<input style="border: none; background-color: #f0f0f0; border-radius: 4px; padding: 2px 10px;" type="text" value="Dynamic"/> ▼

Packet Translation

Advanced Options

ORIGINAL PACKET		TRANSLATED PACKET	
Source Interface	Destination Interface		
<input style="border: none; background-color: #f0f0f0; border-radius: 4px; padding: 2px 10px;" type="text" value="inside"/> ▼	<input style="border: none; background-color: #f0f0f0; border-radius: 4px; padding: 2px 10px;" type="text" value="outside"/>		
Original Address	Original Port	Translated Address	Translated Port
<input style="border: none; background-color: #f0f0f0; border-radius: 4px; padding: 2px 10px;" type="text" value="inside_v6"/> ▼	<input style="border: none; background-color: #f0f0f0; border-radius: 4px; padding: 2px 10px;" type="text" value="Any"/> ▼	<input style="border: none; background-color: #f0f0f0; border-radius: 4px; padding: 2px 10px;" type="text" value="ipv6_pat"/> ▼	<input style="border: none; background-color: #f0f0f0; border-radius: 4px; padding: 2px 10px;" type="text" value="Any"/>

d) Click **OK**.

With this rule, any traffic from the 2001:db8:122:2091::/96 subnet on the inside interface going to the outside interface gets a dynamic PAT66 translation to a port on 2001:db8:122:201b::2.

Monitoring NAT

To monitor and troubleshoot NAT connections, log into the device CLI and use the following commands.

- **show nat** displays the NAT rules and per-rule hit counts. There are additional keywords to show other aspects of NAT.
- **show xlate** displays the actual NAT translations that are currently active.
- **clear xlate** lets you remove an active NAT translation. You might need to remove active translations if you alter NAT rules, because existing connections continue to use the old translation slot until the

connection ends. Clearing a translation allows the system to build a new translation for a client on the client's next connection attempt based on your new rules.

Examples for NAT

The following topics provide examples for configuring NAT on Threat Defense devices.


Providing Access to an Inside Web Server (Static Auto NAT)

The following example performs static NAT for an inside web server. The real address is on a private network, so a public address is required. Static NAT is necessary so hosts can initiate traffic to the web server at a fixed address.



Note This example assumes that the inside interface is not a bridge group interface (BVI) but a standard routed interface. If the inside interface is a BVI, select the specific bridge group member interface to which the web server is attached, for example, inside1_3.

Figure 17: Static NAT for an Inside Web Server



Procedure

- Step 1** Create the network objects that define the server's private and public host addresses.
- Choose **Objects**.
 - Select **Network** from the table of contents and click +.

- c) Define the web server's private address.

Name the network object (for example, WebServerPrivate), select **Host**, and enter the real host IP address, 10.1.2.27.

The screenshot shows the 'New Network Object' configuration window. The title bar is blue with the text 'New Network Object'. Below the title bar, there are four sections: 'Name' with a text input field containing 'WebServerPrivate'; 'Description' with an empty text area; 'Type' with two radio buttons, 'Network' (unselected) and 'Host' (selected); and 'Host' with a text input field containing '10.1.2.27'.

- d) Click **OK**.
e) Click + and define the public address.

Name the network object (for example, WebServerPublic), select **Host**, and enter the host address 209.165.201.10.

The screenshot shows the 'New Network Object' configuration window. The title bar is blue with the text 'New Network Object'. Below the title bar, there are four sections: 'Name' with a text input field containing 'WebServerPublic'; 'Description' with an empty text area; 'Type' with two radio buttons, 'Network' (unselected) and 'Host' (selected); and 'Host' with a text input field containing '209.165.201.10'.

- f) Click **OK**.

Step 2 Configure static NAT for the object.

- a) Select **Policies > NAT**.

- b) Click the + button.
- c) Configure the following properties:
 - **Title** = WebServer (or another name of your choosing).
 - **Create Rule For** = Auto NAT.
 - **Type** = Static.
 - **Source Interface** = inside.
 - **Destination Interface** = outside.
 - **Original Address** = WebServerPrivate network object.
 - **Translated Address** = WebServerPublic network object.

- d) Click **OK**.

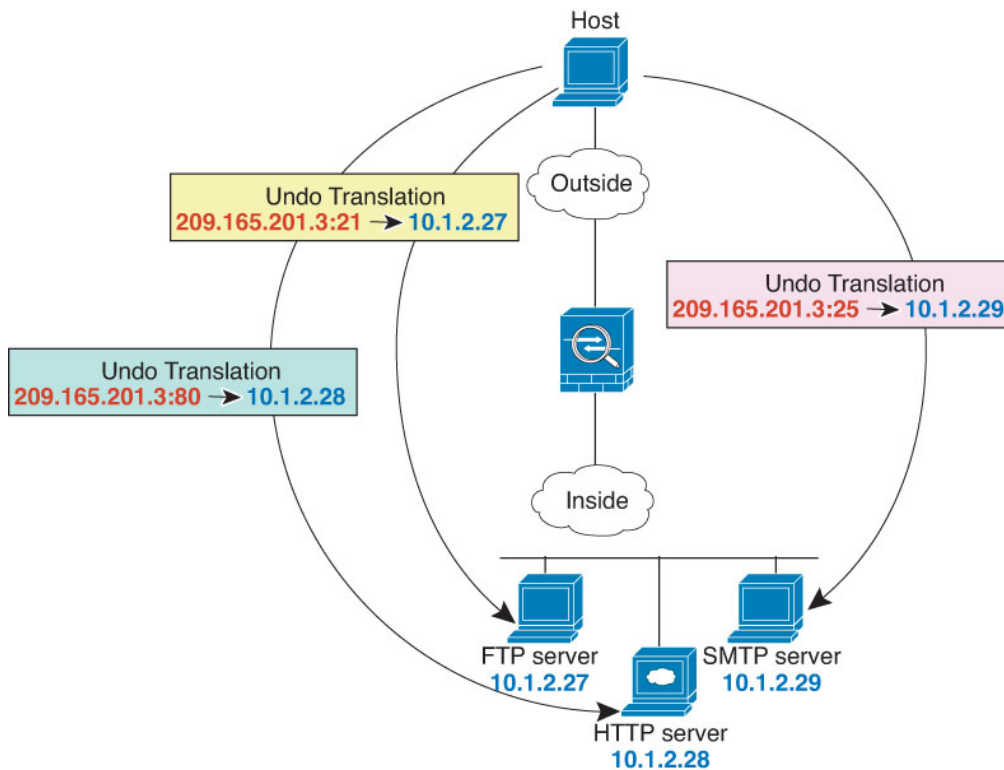
Single Address for FTP, HTTP, and SMTP (Static Auto NAT-with-Port-Translation)

The following static NAT-with-port-translation example provides a single address for remote users to access FTP, HTTP, and SMTP. These servers are actually different devices on the real network, but for each server, you can specify static NAT-with-port-translation rules that use the same mapped IP address, but different ports.



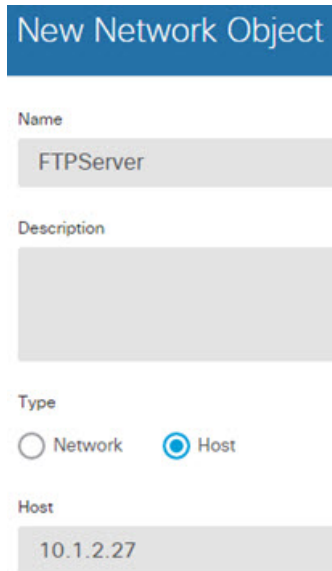
Note This example assumes that the inside interface is a standard routed interface attached to a switch, with the servers attached to the switch. If your inside interface is a bridge group interface (BVI), and the servers are attached to separate bridge group member interfaces, select the specific member interface to which each server is attached for the corresponding rule. For example, the rules might have inside1_2, inside1_3, and inside1_4 for the source interface rather than inside.

Figure 18: Static NAT-with-Port-Translation



Procedure

- Step 1** Create a network object for the FTP server.
- Choose **Objects**.
 - Select **Network** from the table of contents and click +.
 - Name the network object (for example, FTPserver), select **Host**, and enter the real IP address for the FTP server, 10.1.2.27.



New Network Object

Name
FTPServer

Description

Type
 Network Host

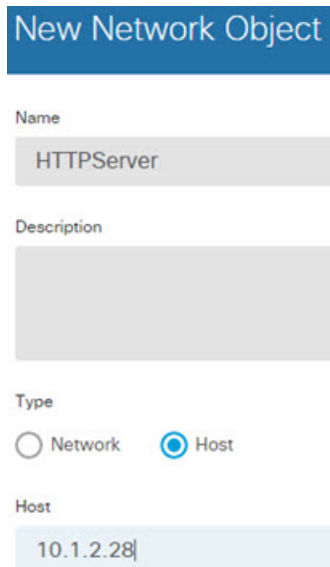
Host
10.1.2.27

d) Click **OK**.

Step 2 Create a network object for the HTTP server.

a) Click +.

b) Name the network object (for example, HTTPserver), select **Host**, and enter the host address 10.1.2.28.



New Network Object

Name
HTTPServer

Description

Type
 Network Host

Host
10.1.2.28

c) Click **OK**.

Step 3 Create a network object for the SMTP server.

a) Click +.

b) Name the network object (for example, SMTPserver), select **Host**, and enter the host address 10.1.2.29.

New Network Object

Name
SMTPServer

Description

Type
 Network Host

Host
10.1.2.29

c) Click **OK**.

Step 4 Create a network object for the public IP address used for the three servers.

- Click +.
- Name the network object (for example, ServerPublicIP), select **Host**, and enter the host address 209.165.201.3.

New Network Object

Name
ServerPublicIP

Description

Type
 Network Host

Host
209.165.201.3

c) Click **OK**.

Step 5 Configure static NAT with port translation for the FTP server, mapping the FTP port to itself.

- Select **Policies > NAT**.
- Click the + button.
- Configure the following properties:

- **Title** = FTPServer (or another name of your choosing).
- **Create Rule For** = Auto NAT.
- **Type** = Static.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Address** = FTPserver network object.
- **Translated Address** = ServerPublicIP network object.
- **Original Port** = FTP port object.
- **Translated Port** = FTP port object.

Add NAT Rule

Title: FTPServer

Create Rule for: Auto NAT

Auto NAT rules translate a specified host or network address regardless of its appearance as the source or destination address of a packet. These rules are automatically ordered and placed in the Auto NAT section.

Placement: Automatically placed in Auto NAT rules

Type: Static

Packet Translation | Advanced Options

Original Packet		Translated Packet	
Source Interface	inside	Destination Interface	outside
Original Address	FTPServer	Translated Address	ServerPublicIP
Original Port	FTP	Translated Port	FTP

d) Click **OK**.

Step 6

Configure static NAT with port translation for the HTTP server, mapping the HTTP port to itself.

- Click the + button.
- Configure the following properties:
 - **Title** = HTTPServer (or another name of your choosing).
 - **Create Rule For** = Auto NAT.
 - **Type** = Static.
 - **Source Interface** = inside.
 - **Destination Interface** = outside.

- **Original Address** = HTTPserver network object.
- **Translated Address** = ServerPublicIP network object.
- **Original Port** = HTTP port object.
- **Translated Port** = HTTP port object.

Add NAT Rule

Title: HTTPServer

Create Rule for: Auto NAT

Auto NAT rules translate a specified host or network address regardless of its appearance as the source or destination address of a packet. These rules are automatically ordered and placed in the Auto NAT section.

Placement: Automatically placed in Auto NAT rules

Type: Static

Packet Translation | Advanced Options

Original Packet

Source Interface: inside

Original Address: HTTPServer

Original Port: HTTP

Translated Packet

Destination Interface: outside

Translated Address: ServerPublicIP

Translated Port: HTTP

c) Click **OK**.

Step 7 Configure static NAT with port translation for the SMTP server, mapping the SMTP port to itself.

- Click the + button.
- Configure the following properties:
 - **Title** = SMTPServer (or another name of your choosing).
 - **Create Rule For** = Auto NAT.
 - **Type** = Static.
 - **Source Interface** = inside.
 - **Destination Interface** = outside.
 - **Original Address** = SMTPserver network object.
 - **Translated Address** = ServerPublicIP network object.
 - **Original Port** = SMTP port object.
 - **Translated Port** = SMTP port object.

Add NAT Rule ?

Title Create Rule for

SMTPServer Auto NAT v

Auto NAT rules translate a specified host or network address regardless of its appearance as the source or destination address of packet. These rules are automatically ordered and placed in the Auto NAT section.

Placement Type

Automatically placed in Auto NAT rules Static v

Packet Translation Advanced Options

Original Packet		Translated Packet	
Source Interface	Destination Interface		
inside v	outside		
Original Address	Original Port	Translated Address	Translated Port
SMTPServer v	SMTP v	ServerPublicIP v	SMTP

c) Click **OK**.


Different Translation Depending on the Destination (Dynamic Manual PAT)

The following figure shows a host on the 10.1.2.0/24 network accessing two different servers. When the host accesses the server at 209.165.201.11, the real address is translated to 209.165.202.129:port. When the host accesses the server at 209.165.200.225, the real address is translated to 209.165.202.130:port.



Note This example assumes that the inside interface is a standard routed interface attached to a switch, with the servers attached to the switch. If your inside interface is a bridge group interface (BVI), and the servers are attached to separate bridge group member interfaces, select the specific member interface to which each server is attached for the corresponding rule. For example, the rules might have inside1_2 and inside1_3 for the source interface rather than inside.

Figure 19: Manual NAT with Different Destination Addresses



Procedure

- Step 1** Create a network object for the inside network.
- Choose **Objects**.
 - Select **Network** from the table of contents and click +.
 - Name the network object (for example, myInsideNetwork), select **Network**, and enter the real network address, 10.1.2.0/24.

New Network Object

Name

myInsideNetwork

Description

Type

 Network
 Host

Network

10.1.2.0/24

d) Click **OK**.

Step 2

Create a network object for the DMZ network 1.

a) Click +.

b) Name the network object (for example, DMZnetwork1), select **Network**, and enter the network address 209.165.201.0/27 (subnet mask of 255.255.255.224).

New Network Object

Name
DMZnetwork1

Description

Type
 Network Host

Network
209.165.201.0/27

c) Click **OK**.

Step 3

Create a network object for the PAT address for DMZ network 1.

a) Click +.

b) Name the network object (for example, PATaddress1), select **Host**, and enter the host address 209.165.202.129.

New Network Object

Name
PATaddress1

Description

Type
 Network Host

Host
209.165.202.129

c) Click **OK**.

- Step 4** Create a network object for the DMZ network 2.
- Click +.
 - Name the network object (for example, DMZnetwork2), select **Network**, and enter the network address 209.165.200.224/27 (subnet mask of 255.255.255.224).

New Network Object

Name
DMZnetwork2

Description

Type
 Network Host

Network
209.165.200.224/27

- Click **OK**.

- Step 5** Create a network object for the PAT address for DMZ network 2.

- Click +.
- Name the network object (for example, PATaddress2), select **Host**, and enter the host address 209.165.202.130.

New Network Object

Name
PATaddress2

Description

Type
 Network Host

Host
209.165.202.130

c) Click **OK**.

Step 6 Configure dynamic manual PAT for DMZ network 1.

a) Select **Policies > NAT**.

b) Click the + button.

c) Configure the following properties:

- **Title** = DMZNetwork1 (or another name of your choosing).
- **Create Rule For** = Manual NAT.
- **Type** = Dynamic.
- **Source Interface** = inside.
- **Destination Interface** = dmz.
- **Original Source Address** = myInsideNetwork network object.
- **Translated Source Address** = PATaddress1 network object.
- **Original Destination Address** = DMZnetwork1 network object.
- **Translated Destination Address** = DMZnetwork1 network object.

Note Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the original and translated destination addresses. Leave all of the port fields blank.

Add NAT Rule

Title: DMZNetwork1 Create Rule for: Manual NAT

Manual NAT rules allow the translation of the source as well as the destination address of a network packet. Destination and port translation are optional. You can place manual NAT rules either before or after Auto NAT rules and insert the rules at a specific location.

Placement: Before Auto NAT Rules Type: Dynamic

Packet Translation Advanced Options

Original Packet		Translated Packet	
Source Interface	inside	Destination Interface	dmz
Source Address	myInsideNetwork	Source Address	PATaddress1
Source Port	Any	Source Port	Any
Destination Address	DMZnetwork1	Destination Address	DMZnetwork1
Destination Port	Any	Destination Port	Any

d) Click **OK**.

Step 7 Configure dynamic manual PAT for DMZ network 2.

a) Click the + button.

b) Configure the following properties:

- **Title** = DMZNetwork2 (or another name of your choosing).
- **Create Rule For** = Manual NAT.
- **Type** = Dynamic.
- **Source Interface** = inside.
- **Destination Interface** = dmz.
- **Original Source Address** = myInsideNetwork network object.
- **Translated Source Address** = PATaddress2 network object.
- **Original Destination Address** = DMZnetwork2 network object.
- **Translated Destination Address** = DMZnetwork2 network object.

Add NAT Rule

Title: DMZNetwork2

Create Rule for: Manual NAT

Manual NAT rules allow the translation of the source as well as the destination address of a network packet. Destination and port translation are optional. You can place manual NAT rules either before or after Auto NAT rules and insert the rules at a specific location.

Placement: Before Auto NAT Rules

Type: Dynamic

Packet Translation | Advanced Options

Original Packet		Translated Packet	
Source Interface	inside	Destination Interface	dmz
Source Address	myInsideNetwork	Source Address	PATaddress2
Source Port	Any	Source Port	Any
Destination Address	DMZnetwork2	Destination Address	DMZnetwork2
Destination Port	Any	Destination Port	Any

c) Click **OK**.

Different Translation Depending on the Destination Address and Port (Dynamic Manual PAT)


The following figure shows the use of source and destination ports. The host on the 10.1.2.0/24 network accesses a single host for both web services and Telnet services. When the host accesses the server for Telnet services, the real address is translated to 209.165.202.129:port. When the host accesses the same server for web services, the real address is translated to 209.165.202.130:port.



Note

This example assumes that the inside interface is a standard routed interface attached to a switch, with the server attached to the switch. If your inside interface is a bridge group interface (BVI), and the server is attached to a bridge group member interface, select the specific member interface to which the server is attached. For example, the rule might have inside1_2 for the source interface rather than inside.

Figure 20: Manual NAT with Different Destination Ports



Procedure

Step 1

Create a network object for the inside network.

- Choose **Objects**.
- Select **Network** from the table of contents and click +.
- Name the network object (for example, myInsideNetwork), select **Network**, and enter the real network address, 10.1.2.0/24.

New Network Object

Name

myInsideNetwork

Description

Type

Network Host

Network

10.1.2.0/24

d) Click **OK**.

Step 2 Create a network object for the Telnet/Web server.

a) Click +.

b) Name the network object (for example, TelnetWebServer), select **Host**, and enter the host address 209.165.201.11.

New Network Object

Name
TelnetWebServer

Description

Type
 Network Host

Host
209.165.201.11

c) Click **OK**.

Step 3 Create a network object for the PAT address when using Telnet.

a) Click +.

b) Name the network object (for example, PATAddress1), select **Host**, and enter the host address 209.165.202.129.

New Network Object

Name
PATAddress1

Description

Type
 Network Host

Host
209.165.202.129

c) Click **OK**.

Step 4 Create a network object for the PAT address when using HTTP.

- a) Click +.
- b) Name the network object (for example, PATaddress2), select **Host**, and enter the host address 209.165.202.130.

New Network Object

Name
PATaddress2

Description

Type
 Network Host

Host
209.165.202.130

- c) Click **OK**.

Step 5 Configure dynamic manual PAT for Telnet access.

- a) Select **Policies > NAT**.
- b) Click the + button.
- c) Configure the following properties:
 - **Title** = TelnetServer (or another name of your choosing).
 - **Create Rule For** = Manual NAT.
 - **Type** = Dynamic.
 - **Source Interface** = inside.
 - **Destination Interface** = dmz.
 - **Original Source Address** = myInsideNetwork network object.
 - **Translated Source Address** = PATaddress1 network object.
 - **Original Destination Address** = TelnetWebServer network object.
 - **Translated Destination Address** = TelnetWebServer network object.
 - **Original Destination Port** = TELNET port object.
 - **Translated Destination Port** = TELNET port object.

Note Because you do not want to translate the destination address or port, you need to configure identity NAT for them by specifying the same address for the original and translated destination addresses, and the same port for the original and translated port.

Add NAT Rule

Title: TelnetServer

Create Rule for: Manual NAT

Manual NAT rules allow the translation of the source as well as the destination address of a network packet. Destination and port translation are optional. You can place manual NAT rules either before or after Auto NAT rules and insert the rules at a specific location.

Placement: Before Auto NAT Rules

Type: Dynamic

Packet Translation | Advanced Options

Original Packet		Translated Packet	
Source Interface	inside	Destination Interface	dmz
Source Address	myInsideNetwork	Source Address	PATAddress1
Source Port	Any	Source Port	Any
Destination Address	TelnetWebServe	Destination Address	TelnetWebServe
Destination Port	TELNET	Destination Port	TELNET

d) Click **OK**.

Step 6 Configure dynamic manual PAT for web access.

a) Click the + button.

b) Configure the following properties:

- **Title** = WebServer (or another name of your choosing).
- **Create Rule For** = Manual NAT.
- **Type** = Dynamic.
- **Source Interface** = inside.
- **Destination Interface** = dmz.
- **Original Source Address** = myInsideNetwork network object.
- **Translated Source Address** = PATaddress2 network object.
- **Original Destination Address** = TelnetWebServer network object.
- **Translated Destination Address** = TelnetWebServer network object.
- **Original Destination Port** = HTTP port object.
- **Translated Destination Port** = HTTP port object.

Add NAT Rule

Title: WebServer

Create Rule for: Manual NAT

Manual NAT rules allow the translation of the source as well as the destination address of a network packet. Destination and port translation are optional. You can place manual NAT rules either before or after Auto NAT rules and insert the rules at a specific location.

Placement: Before Auto NAT Rules

Type: Dynamic

Packet Translation | Advanced Options

Original Packet		Translated Packet	
Source Interface	inside	Destination Interface	dmz
Source Address	myInsideNetwork	Source Address	PATAddress2
Source Port	Any	Source Port	Any
Destination Address	TelnetWebServe	Destination Address	TelnetWebServe
Destination Port	HTTP	Destination Port	HTTP

c) Click **OK**.

Rewriting DNS Queries and Responses Using NAT

You might need to configure the FTD device to modify DNS replies by replacing the address in the reply with an address that matches the NAT configuration. You can configure DNS modification when you configure each translation rule. DNS modification is also known as DNS doctoring.

This feature rewrites the address in DNS queries and replies that match a NAT rule (for example, the A record for IPv4, the AAAA record for IPv6, or the PTR record for reverse DNS queries). For DNS replies traversing from a mapped interface to any other interface, the record is rewritten from the mapped value to the real value. Inversely, for DNS replies traversing from any interface to a mapped interface, the record is rewritten from the real value to the mapped value. This feature works with NAT44, NAT 66, NAT46, and NAT64.

Following are the main circumstances when you would need to configure DNS rewrite on a NAT rule.

- The rule is NAT64 or NAT46, and the DNS server is on the outside network. You need DNS rewrite to convert between DNS A records (for IPv4) and AAAA records (for IPv6).
- The DNS server is on the outside, clients are on the inside, and some of the fully-qualified domain names that the clients use resolve to other inside hosts.
- The DNS server is on the inside and responds with private IP addresses, clients are on the outside, and the clients access fully-qualified domain names that point to servers that are hosted on the inside.

DNS Rewrite Limitations

Following are some limitations with DNS rewrite:

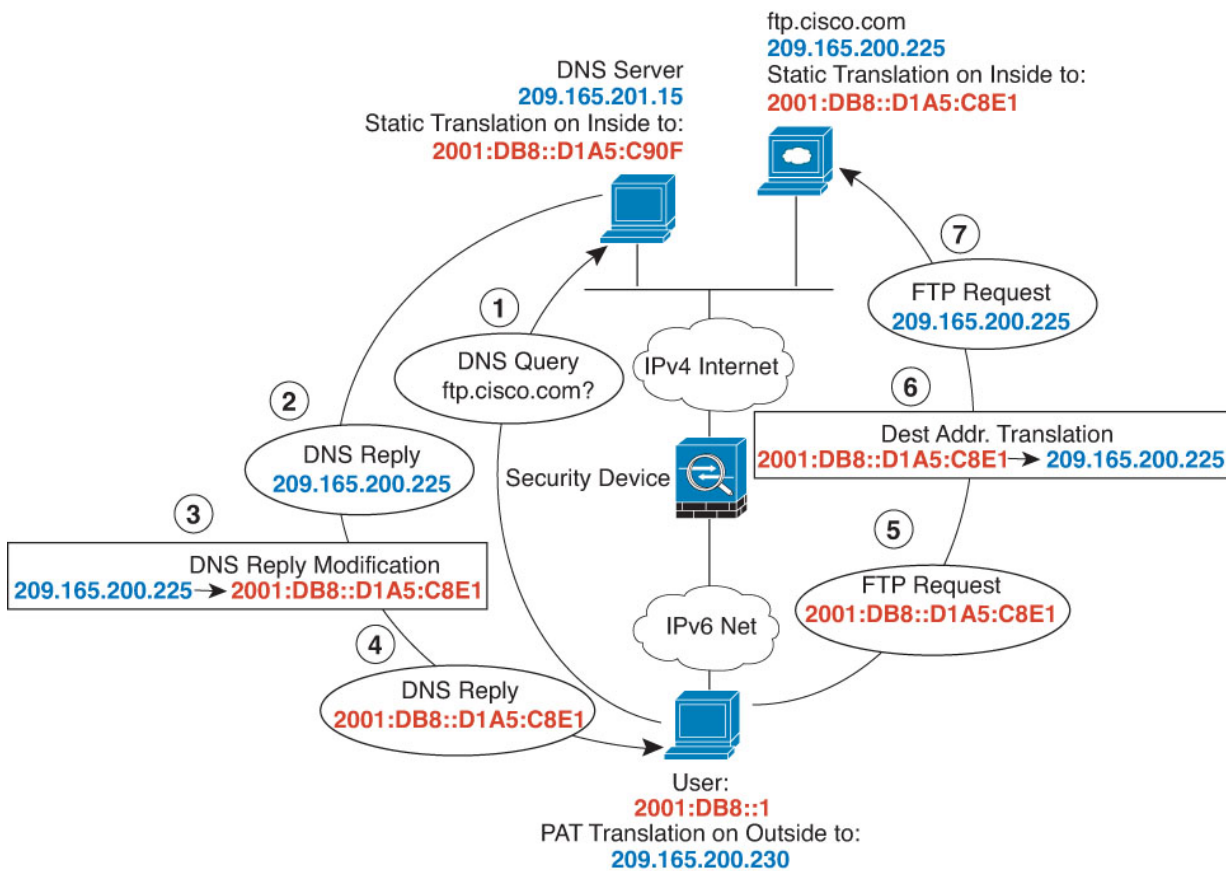
- DNS rewrite is not applicable for PAT because multiple PAT rules are applicable for each A or AAAA record, and the PAT rule to use is ambiguous.
- If you configure a manual NAT rule, you cannot configure DNS modification if you specify the destination address as well as the source address. These kinds of rules can potentially have a different translation for a single address when going to A vs. B. Therefore, they can not accurately match the IP address inside the DNS reply to the correct NAT rule; the DNS reply does not contain information about which source/destination address combination was in the packet that prompted the DNS request.
- DNS rewrite is actually done on the xlate entry, not the NAT rule. Thus, if there is no xlate for a dynamic rule, rewrite cannot be done correctly. The same problem does not occur for static NAT.
- DNS rewrite does not rewrite DNS Dynamic Update messages (opcode 5).

The following topics provide examples of DNS rewrite in NAT rules.

DNS 64 Reply Modification

The following figure shows an FTP server and DNS server on the outside IPv4 network. The system has a static translation for the outside server. In this case, when an inside IPv6 user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.200.225.

Because you want inside users to use the mapped address for ftp.cisco.com (2001:DB8::D1A5:C8E1, where D1A5:C8E1 is the IPv6 equivalent of 209.165.200.225) you need to configure DNS reply modification for the static translation. This example also includes a static NAT translation for the DNS server, and a PAT rule for the inside IPv6 hosts.



Note This example assumes that the inside interface is not a bridge group interface (BVI) but a standard routed interface. If the inside interface is a BVI, you need to duplicate the rules for each member interface.

Procedure

Step 1 Create the network objects for the FTP server, DNS server, inside network, and PAT pool.

- a) Choose **Objects**.
- b) Select **Network** from the table of contents and click +.
- c) Define the real FTP server address.

Name the network object (for example, ftp_server), select **Host**, and enter the real host IP address, 209.165.200.225.

Add Network Object

Name
ftp_server

Description

Type
 Network Host

Host
209.165.200.225

- d) Click **OK**.
- e) Click + and define the DNS server's real address.

Name the network object (for example, dns_server), select **Host**, and enter the host address 209.165.201.15.

Add Network Object

Name
dns_server

Description

Type
 Network Host

Host
209.165.201.15

- f) Click **OK**.
- g) Click + and define the inside IPv6 network.

Name the network object (for example, inside_v6), select **Network**, and enter the network address, 2001:DB8::/96.

Add Network Object

Name
inside_v6

Description

Type
 Network Host

Network
2001:DB8::/96

- h) Click **OK**.
- i) Click + and define the IPv4 PAT address for the inside IPv6 network.
Name the network object (for example, ipv4_pat), select **Host**, and enter the host address, 209.165.200.230.

Add Network Object

Name
ipv4_pat

Description

Type
 Network Host

Host
209.165.200.230

- j) Click **OK**.

Step 2 Configure the static NAT rule with DNS modification for the FTP server.

- Select **Policies > NAT**.
- Click the + button.
- Configure the following properties:
 - **Title** = FTPServer (or another name of your choosing).
 - **Create Rule For** = Auto NAT.

- **Type** = Static.
- **Source Interface** = outside.
- **Destination Interface** = inside.
- **Original Address** = ftp_server network object.
- **Translated Address** = inside_v6 network object. Because the IPv4 embedded address method is used when converting IPv4 to IPv6 addresses, 209.165.200.225 is converted to the IPv6 equivalent D1A5:C8E1 and the network prefix is added to get the full address, 2001:DB8::D1A5:C8E1.
- On the **Advanced Options** tab, select **Translate DNS replies that match this rule**.

Add NAT Rule

Title: FTPServer Create Rule for: Auto NAT Status:

Auto NAT rules translate a specified host or network address regardless of its appearance as the source or destination address of a packet. These rules are automatically ordered and placed in the Auto NAT section.

Placement: Automatically placed in Auto NAT rules Type: Static

Packet Translation Advanced Options

ORIGINAL PACKET

Source Interface: outside

Original Address: ftp_server Original Port: Any

TRANSLATED PACKET

Destination Interface: inside

Translated Address: inside_v6 Translated Port: Any

d) Click **OK**.

Step 3

Configure the static NAT rule for the DNS server.

- Select **Policies > NAT**.
- Click the + button.
- Configure the following properties:
 - **Title** = DNSServer (or another name of your choosing).
 - **Create Rule For** = Auto NAT.
 - **Type** = Static.
 - **Source Interface** = outside.
 - **Destination Interface** = inside.

- **Original Address** = dns_server network object.
- **Translated Address** = inside_v6 network object. Because the IPv4 embedded address method is used when converting IPv4 to IPv6 addresses, 209.165.201.15 is converted to the IPv6 equivalent D1A5:C90F and the network prefix is added to get the full address, 2001:DB8::D1A5:C90F.

Add NAT Rule

Title: DNSServer Create Rule for: Auto NAT Status:

Auto NAT rules translate a specified host or network address regardless of its appearance as the source or destination address of a packet. These rules are automatically ordered and placed in the Auto NAT section.

Placement: Automatically placed in Auto NAT rules Type: Static

Packet Translation Advanced Options

ORIGINAL PACKET		TRANSLATED PACKET	
Source Interface	outside	Destination Interface	inside
Original Address	dns_server	Translated Address	inside_v6
Original Port	Any	Translated Port	Any

d) Click **OK**.

Step 4 Configure the dynamic PAT rule for the inside IPv6 network.

- Select **Policies > NAT**.
- Click the + button.
- Configure the following properties:
 - **Title** = PAT64Rule (or another name of your choosing).
 - **Create Rule For** = Auto NAT.
 - **Type** = Dynamic.
 - **Source Interface** = inside.
 - **Destination Interface** = outside.
 - **Original Address** = inside_v6 network object.
 - **Translated Address** = ipv4_pat network object.

Add NAT Rule ?

Title	Create Rule for	Status
PAT64Rule	Auto NAT v	<input checked="" type="checkbox"/>

Auto NAT rules translate a specified host or network address regardless of its appearance as the source or destination address of a packet. These rules are automatically ordered and placed in the Auto NAT section.

Placement	Type
Automatically placed in Auto NAT rules	Dynamic v

Packet Translation

Advanced Options

ORIGINAL PACKET		TRANSLATED PACKET	
Source Interface	Destination Interface		
inside v	outside		
Original Address	Original Port	Translated Address	Translated Port
inside_v6 v	Any v	ipv4_pat v	Any


d) Click **OK**.

DNS Reply Modification, DNS Server on Outside

The following figure shows a DNS server that is accessible from the outside interface. A server, ftp.cisco.com, is on the inside interface. You configure NAT to statically translate the ftp.cisco.com real address (10.1.3.14) to a mapped address (209.165.201.10) that is visible on the outside network.

In this case, you want to enable DNS reply modification on this static rule so that inside users who have access to ftp.cisco.com using the real address receive the real address from the DNS server, and not the mapped address.

When an inside host sends a DNS request for the address of ftp.cisco.com, the DNS server replies with the mapped address (209.165.201.10). The system refers to the static rule for the inside server and translates the address inside the DNS reply to 10.1.3.14. If you do not enable DNS reply modification, then the inside host attempts to send traffic to 209.165.201.10 instead of accessing ftp.cisco.com directly.



Note This example assumes that the inside interface is not a bridge group interface (BVI) but a standard routed interface. If the inside interface is a BVI, you need to duplicate the rules for each member interface.

Procedure

Step 1 Create the network objects for the FTP server.

- Choose **Objects**.
- Select **Network** from the table of contents and click +.
- Define the real FTP server address.

Name the network object (for example, ftp_server), select **Host**, and enter the real host IP address, 10.1.3.14.

Add Network Object

Name
ftp_server

Description

Type
 Network Host

Host
10.1.3.14

- d) Click **OK**.
- e) Click + and define the FTP server's translated address.

Name the network object (for example, ftp_server_outside), select **Host**, and enter the host address 209.165.201.10.

Add Network Object

Name
ftp_server_outside

Description

Type
 Network Host

Host
209.165.201.10

Step 2 Configure the static NAT rule with DNS modification for the FTP server.

- a) Select **Policies > NAT**.
- b) Click the + button.
- c) Configure the following properties:
 - **Title** = FTPServer (or another name of your choosing).
 - **Create Rule For** = Auto NAT.

- **Type** = Static.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Address** = ftp_server network object.
- **Translated Address** = ftp_server_outside network object.
- On the **Advanced Options** tab, select **Translate DNS replies that match this rule**.

Add NAT Rule ?

Title	Create Rule for	Status
FTPServer	Auto NAT ▼	<input checked="" type="checkbox"/>

Auto NAT rules translate a specified host or network address regardless of its appearance as the source or destination address of a packet. These rules are automatically ordered and placed in the Auto NAT section.

Placement	Type
Automatically placed in Auto NAT rules	Static ▼


Packet Translation Advanced Options

ORIGINAL PACKET		TRANSLATED PACKET	
Source Interface		Destination Interface	
inside ▼		outside	
Original Address	Original Port	Translated Address	Translated Port
ftp_server ▼	Any ▼	ftp_server_outside ▼	Any

d) Click **OK**.

DNS Reply Modification, DNS Server on Host Network

The following figure shows an FTP server and DNS server on the outside. The system has a static translation for the outside server. In this case, when an inside user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.20.10. Because you want inside users to use the mapped address for ftp.cisco.com (10.1.2.56) you need to configure DNS reply modification for the static translation.



Note This example assumes that the inside interface is not a bridge group interface (BVI) but a standard routed interface. If the inside interface is a BVI, you need to duplicate the rules for each member interface.

Procedure

Step 1 Create the network objects for the FTP server.

- Choose **Objects**.
- Select **Network** from the table of contents and click +.
- Define the real FTP server address.

Name the network object (for example, ftp_server), select **Host**, and enter the real host IP address, 209.165.201.10.

Add Network Object

Name
ftp_server

Description

Type
 Network Host

Host
209.165.201.10

- d) Click **OK**.
- e) Click + and define the FTP server's translated address.

Name the network object (for example, ftp_server_translated), select **Host**, and enter the host address 10.1.2.56.

Add Network Object

Name
ftp_server_translated

Description

Type
 Network Host

Host
10.1.2.56

Step 2 Configure the static NAT rule with DNS modification for the FTP server.

- a) Select **Policies > NAT**.
- b) Click the + button.
- c) Configure the following properties:
 - **Title** = FTPServer (or another name of your choosing).
 - **Create Rule For** = Auto NAT.

- **Type** = Static.
- **Source Interface** = outside.
- **Destination Interface** = inside.
- **Original Address** = ftp_server network object.
- **Translated Address** = ftp_server_translated network object.
- On the **Advanced Options** tab, select **Translate DNS replies that match this rule**.

Add NAT Rule

Title: FTPServer Create Rule for: Auto NAT Status:

Auto NAT rules translate a specified host or network address regardless of its appearance as the source or destination address of a packet. These rules are automatically ordered and placed in the Auto NAT section.

Placement: Automatically placed in Auto NAT rules Type: Static

Packet Translation Advanced Options

ORIGINAL PACKET		TRANSLATED PACKET	
Source Interface	outside	Destination Interface	inside
Original Address	ftp_server	Translated Address	ftp_server_transla
Original Port	Any	Translated Port	Any

d) Click **OK**.



PART **III**

Virtual Private Networks (VPN)

- [Site-to-Site VPN, on page 249](#)



CHAPTER 11

Site-to-Site VPN

A virtual private network (VPN) is a network connection that establishes a secure tunnel between remote peers using a public source, such as the Internet or other network. VPNs use tunnels to encapsulate data packets within normal IP packets for forwarding over IP-based networks. They use encryption to ensure privacy and authentication to ensure the integrity of data.

- [VPN Basics, on page 249](#)
- [Managing Site-to-Site VPNs, on page 253](#)
- [Monitoring Site-to-Site VPN, on page 266](#)
- [Examples for Site-to-Site VPN, on page 266](#)

VPN Basics

Tunneling makes it possible to use a public TCP/IP network, such as the Internet, to create secure connections between remote users and private corporate networks. Each secure connection is called a tunnel.

IPsec-based VPN technologies use the Internet Security Association and Key Management Protocol (ISAKMP, or IKE) and IPsec tunneling standards to build and manage tunnels. ISAKMP and IPsec accomplish the following:

- Negotiate tunnel parameters.
- Establish tunnels.
- Authenticate users and data.
- Manage security keys.
- Encrypt and decrypt data.
- Manage data transfer across the tunnel.
- Manage data transfer inbound and outbound as a tunnel endpoint or router.

A device in a VPN functions as a bidirectional tunnel endpoint. It can receive plain packets from the private network, encapsulate them, create a tunnel, and send them to the other end of the tunnel where they are unencapsulated and sent to their final destination. It can also receive encapsulated packets from the public network, unencapsulate them, and send them to their final destination on the private network.

After the site-to-site VPN connection is established, the hosts behind the local gateway can connect to the hosts behind the remote gateway through the secure VPN tunnel. A connection consists of the IP addresses

and hostnames of the two gateways, the subnets behind them, and the method the two gateways use to authenticate to each other.

Internet Key Exchange (IKE)

Internet Key Exchange (IKE) is a key management protocol that is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and to automatically establish IPsec security associations (SAs).

The IKE negotiation comprises two phases. Phase 1 negotiates a security association between two IKE peers, which enables the peers to communicate securely in Phase 2. During Phase 2 negotiation, IKE establishes SAs for other applications, such as IPsec. Both phases use proposals when they negotiate a connection.

An IKE policy is a set of algorithms that two peers use to secure the IKE negotiation between them. IKE negotiation begins by each peer agreeing on a common (shared) IKE policy. This policy states which security parameters protect subsequent IKE negotiations. For IKE version 1 (IKEv1), IKE policies contain a single set of algorithms and a modulus group. Unlike IKEv1, in an IKEv2 policy, you can select multiple algorithms and modulus groups from which peers can choose during the Phase 1 negotiation. It is possible to create a single IKE policy, although you might want different policies to give higher priority to your most desired options. For site-to-site VPNs, you can create a single IKE policy.

To define an IKE policy, specify:

- A unique priority (1 to 65,543, with 1 the highest priority).
- An encryption method for the IKE negotiation, to protect the data and ensure privacy.
- A Hashed Message Authentication Codes (HMAC) method (called integrity algorithm in IKEv2) to ensure the identity of the sender, and to ensure that the message has not been modified in transit.
- For IKEv2, a separate pseudorandom function (PRF) used as the algorithm to derive keying material and hashing operations required for the IKEv2 tunnel encryption. The options are the same as those used for the hash algorithm.
- A Diffie-Hellman group to determine the strength of the encryption-key-determination algorithm. The device uses this algorithm to derive the encryption and hash keys.
- An authentication method, to ensure the identity of the peers.
- A limit to the time the device uses an encryption key before replacing it.

When IKE negotiation begins, the peer that starts the negotiation sends all of its enabled policies to the remote peer, and the remote peer searches for a match with its own policies, in priority order. A match between IKE policies exists if they have the same encryption, hash (integrity and PRF for IKEv2), authentication, and Diffie-Hellman values, and an SA lifetime less than or equal to the lifetime in the policy sent. If the lifetimes are not identical, the shorter lifetime, obtained from the remote peer, applies. By default, a simple IKE policy that uses DES is the only enabled policy. You can enable other IKE policies at higher priorities to negotiate stronger encryption standards, but the DES policy should ensure a successful negotiation.

How Secure Should a VPN Connection Be?

Because a VPN tunnel typically traverses a public network, most likely the Internet, you need to encrypt the connection to protect the traffic. You define the encryption and other security techniques to apply using IKE policies and IPsec proposals.

If your device license allows you to apply strong encryption, there is a wide range of encryption and hash algorithms, and Diffie-Hellman groups, from which to choose. However, as a general rule, the stronger the encryption that you apply to the tunnel, the worse the system performance. Find a balance between security and performance that provides sufficient protection without compromising efficiency.

We cannot provide specific guidance on which options to choose. If you operate within a larger corporation or other organization, there might already be defined standards that you need to meet. If not, take the time to research the options.

The following topics explain the available options.

Deciding Which Encryption Algorithm to Use

When deciding which encryption algorithms to use for the IKE policy or IPsec proposal, your choice is limited to algorithms supported by the devices in the VPN.

For IKEv2, you can configure multiple encryption algorithms. The system orders the settings from the most secure to the least secure and negotiates with the peer using that order. For IKEv1, you can select a single option only.

For IPsec proposals, the algorithm is used by the Encapsulating Security Protocol (ESP), which provides authentication, encryption, and anti-replay services. ESP is IP protocol type 50. In IKEv1 IPsec proposals, the algorithm name is prefixed with ESP-.

If your device license qualifies for strong encryption, you can choose from the following encryption algorithms. If you are not qualified for strong encryption, you can select DES only.

- **AES-GCM**—(IKEv2 only.) Advanced Encryption Standard in Galois/Counter Mode is a block cipher mode of operation providing confidentiality and data-origin authentication, and provides greater security than AES. AES-GCM offers three different key strengths: 128-, 192-, and 256-bit keys. A longer key provides higher security but a reduction in performance. GCM is a mode of AES that is required to support NSA Suite B. NSA Suite B is a set of cryptographic algorithms that devices must support to meet federal standards for cryptographic strength.
- **AES-GMAC**—(IKEv2 IPsec proposals only.) Advanced Encryption Standard Galois Message Authentication Code is a block cipher mode of operation providing only data-origin authentication. It is a variant of AES-GCM that allows data authentication without encrypting the data. AES-GMAC offers three different key strengths: 128-, 192-, and 256-bit keys.
- **AES**—Advanced Encryption Standard is a symmetric cipher algorithm that provides greater security than DES and is computationally more efficient than 3DES. AES offers three different key strengths: 128-, 192-, and 256-bit keys. A longer key provides higher security but a reduction in performance.
- **3DES**—Triple DES, which encrypts three times using 56-bit keys, is more secure than DES because it processes each block of data three times with a different key. However, it uses more system resources and is slower than DES.
- **DES**—Data Encryption Standard, which encrypts using 56-bit keys, is a symmetric secret-key block algorithm. If your license account does not meet the requirements for export controls, this is your only option. It is faster than 3DES and uses less system resources, but it is also less secure. If you do not need strong data confidentiality, and if system resources or speed is a concern, choose DES.
- **Null, ESP-Null**—Do not use. A null encryption algorithm provides authentication without encryption. This is not supported on most platforms.

Deciding Which Hash Algorithms to Use

In IKE policies, the hash algorithm creates a message digest, which is used to ensure message integrity. In IKEv2, the hash algorithm is separated into two options, one for the integrity algorithm, and one for the pseudo-random function (PRF).

In IPsec proposals, the hash algorithm is used by the Encapsulating Security Protocol (ESP) for authentication. In IKEv2 IPsec Proposals, this is called the integrity hash. In IKEv1 IPsec proposals, the algorithm name is prefixed with ESP-, and there is also an -HMAC suffix (which stands for “hash method authentication code”).

For IKEv2, you can configure multiple hash algorithms. The system orders the settings from the most secure to the least secure and negotiates with the peer using that order. For IKEv1, you can select a single option only.

You can choose from the following hash algorithms.

- **SHA (Secure Hash Algorithm)**—Standard SHA (SHA1) produces a 160-bit digest. SHA is more resistant to brute-force attacks than MD5. However, it is also more resource intensive than MD5. For implementations that require the highest level of security, use the SHA hash algorithm.

The following SHA-2 options, which are even more secure, are available for IKEv2 configurations. Choose one of these if you want to implement the NSA Suite B cryptography specification.

- **SHA256**—Specifies the Secure Hash Algorithm SHA 2 with the 256-bit digest.
- **SHA384**—Specifies the Secure Hash Algorithm SHA 2 with the 384-bit digest.
- **SHA512**—Specifies the Secure Hash Algorithm SHA 2 with the 512-bit digest.
- **MD5 (Message Digest 5)**—Produces a 128-bit digest. MD5 uses less processing time for an overall faster performance than SHA, but it is considered to be weaker than SHA.
- **Null or None (NULL, ESP-NONE)**—(IPsec Proposals only.) A null Hash Algorithm; this is typically used for testing purposes only. However, you should choose the null integrity algorithm if you select one of the AES-GCM/GMAC options as the encryption algorithm. Even if you choose a non-null option, the integrity hash is ignored for these encryption standards.

Deciding Which Diffie-Hellman Modulus Group to Use

You can use the following Diffie-Hellman key derivation algorithms to generate IPsec security association (SA) keys. Each group has a different size modulus. A larger modulus provides higher security, but requires more processing time. You must have a matching modulus group on both peers.

If you select AES encryption, to support the large key sizes required by AES, you should use Diffie-Hellman (DH) Group 5 or higher. IKEv1 policies do not support all of the groups listed below.

To implement the NSA Suite B cryptography specification, use IKEv2 and select one of the elliptic curve Diffie-Hellman (ECDH) options: 19, 20, or 21. Elliptic curve options and groups that use 2048-bit modulus are less exposed to attacks such as Logjam.

For IKEv2, you can configure multiple groups. The system orders the settings from the most secure to the least secure and negotiates with the peer using that order. For IKEv1, you can select a single option only.


- **1**—Diffie-Hellman Group 1: 768-bit modulus. DH group 1 is considered insecure, please do not use it.
- **2**—Diffie-Hellman Group 2: 1024-bit modular exponential (MODP) group. This option is no longer considered good protection.

- 5—Diffie-Hellman Group 5: 1536-bit MODP group. Formerly considered good protection for 128-bit keys, this option is no longer considered good protection.
- 14—Diffie-Hellman Group 14: 2048-bit modular exponential (MODP) group. Considered good protection for 192-bit keys.
- 19—Diffie-Hellman Group 19: National Institute of Standards and Technology (NIST) 256-bit elliptic curve modulo a prime (ECP) group.
- 20—Diffie-Hellman Group 20: NIST 384-bit ECP group.
- 21—Diffie-Hellman Group 21: NIST 521-bit ECP group.
- 24—Diffie-Hellman Group 24: 2048-bit MODP group with 256-bit prime order subgroup. This option is no longer recommended.

VPN Topologies

You can configure only point-to-point VPN connections using FDM. Although all connections are point-to-point, you can link into larger hub-and-spoke or meshed VPNs by defining each of the tunnels in which your device participates.

The following diagram displays a typical point-to-point VPN topology. In a point-to-point VPN topology, two endpoints communicate directly with each other. You configure the two endpoints as peer devices, and either device can start the secured connection.



Managing Site-to-Site VPNs

A virtual private network (VPN) is a network connection that establishes a secure tunnel between remote peers using a public source, such as the Internet or other network. VPNs use tunnels to encapsulate data packets within normal IP packets for forwarding over IP-based networks. They use encryption to ensure privacy and authentication to ensure the integrity of data.

You can create VPN connections to peer devices. All connections are point-to-point, but you can link the device into larger hub-and-spoke or meshed VPNs by configuring all relevant connections.

Before you begin

The following facts control the type and number of site-to-site VPN connections that you can recreate:

- VPN connections use encryption to secure network privacy. The encryption algorithms that you can use depend on whether your base license allows strong encryption. This is controlled by whether you selected the option to allow export-controlled functionality on the device when you registered with Cisco Smart

License Manager. If you are using the evaluation license, or you did not enable export-controlled functionality, you cannot use strong encryption.




- You can create at most 20 unique IPsec profiles. Uniqueness is determined by the combination of IKEv1/v2 proposals and certificates, connection type, DH group and SA lifetime. You can reuse existing profiles. Thus, if you use the same settings for all your site-to-site VPN connections, you have one unique IPsec profile. Once you reach the limit of 20 unique IPsec profiles, you cannot create new site-to-site VPN connections unless you use the same combination of attributes that you used for an existing connection profile.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Site-to-Site VPN group.

This opens the Site-to-Site VPN page, which lists all of the connections that you have configured.

Step 2 Do any of the following.

- To create a new Site-to-Site VPN connection, click the + button. See [Configuring a Site-to-Site VPN Connection, on page 254](#).
If there are no connections yet, you can also click the **Create Site-to-Site Connection** button.
- To edit an existing connection, click the edit icon () for the connection. See [Configuring a Site-to-Site VPN Connection, on page 254](#).
- To copy a summary of the connection configuration to the clipboard, click the copy icon () for the connection. You can paste this information in a document and send it to the administrator for the remote device to help configure that end of the connection.
- To delete a connection that you no longer need, click the delete icon () for the connection.

Configuring a Site-to-Site VPN Connection

You can create a point-to-point VPN connection to link your device to another device, assuming that you have the cooperation and permission of the remote device owner. Although all connections are point-to-point, you can link into larger hub-and-spoke or meshed VPNs by defining each of the tunnels in which your device participates.





Note You can create a single VPN connection per local network/remote network combination. However, you can create multiple connections for a local network if the remote network is unique in each connection profile.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Site-to-Site VPN group.

Step 2 Do any of the following:

- To create a new Site-to-Site VPN connection, click the + button.
If there are no connections yet, you can also click the **Create Site-to-Site Connection** button.
- To edit an existing connection, click the edit icon () for the connection.

To delete a connection that you no longer need, click the delete icon () for the connection.

Step 3 Define the endpoints of the point-to-point VPN connection.

- **Connection Profile Name**—The name for this connection, up to 64 characters without spaces. For example, MainOffice. You cannot use an IP address as the name.
- **Local Site**—These options define the local endpoint.
 - **Local VPN Access Interface**—Select the interface to which the remote peer can connect. This is typically the outside interface. The interface cannot be a member of a bridge group.
 - **Local Network**—Click + and select the network objects that identify the local networks that should participate in the VPN connection. Users on these networks will be able to reach the remote networks through the connection.

Note You can use IPv4 or IPv6 addresses for these networks, but you must have a matching address type on each side of the connection. For example, the VPN connection for a local IPv4 network must have at least one remote IPv4 network. You can combine IPv4 and IPv6 on both sides of a single connection. The protected networks for the endpoints cannot overlap.

- **Remote Site**—These options define the remote endpoint.
 - **Remote IP Address**—Enter the IP address of the remote VPN peer's interface that will host the VPN connection.
 - **Remote Network**—Click + and select the network objects that identify the remote networks that should participate in the VPN connection. Users on these networks will be able to reach the local networks through the connection.

Step 4 Click **Next**.

Step 5 Define the privacy configuration for the VPN.

Note Your license determines which encryption protocols you can select. You must qualify for strong encryption, i.e. satisfy export controls, to choose any but the most basic options.

- **IKE Version 2, IKE Version 1**—Choose the IKE versions to use during Internet Key Exchange (IKE) negotiations. Select either or both options as appropriate. When the device attempts to negotiate a connection with the other peer, it uses whichever versions you allow and that the other peer accepts. If you allow both versions, the device automatically falls back to the other version if negotiations are unsuccessful with the initially chosen version. IKEv2 is always tried first if it is configured. Both peers must support IKEv2 to use it in a negotiation.

- **IKE Policy**—Internet Key Exchange (IKE) is a key management protocol that is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and automatically establish IPsec security associations (SAs). This is a global policy: the objects you enable are applied to all VPNs. Click **Edit** to examine the current globally-enabled policies per IKE version, and to enable and create new policies. For more information, see [Configuring the Global IKE Policy, on page 256](#).
- **IPsec Proposal**—The IPsec proposal defines the combination of security protocols and algorithms that secure traffic in an IPsec tunnel. Click **Edit** and select the proposals for each IKE version. Select all proposals that you want to allow. Click **Set Default** to simply select the system defaults, which differ based on your export compliance. The system negotiates with the peer, starting from the strongest to the weakest proposal, until a match is agreed upon. For more information, see [Configuring IPsec Proposals, on page 261](#).
- **(IKEv2) Local Preshared Key, Remote Peer Preshared Key**—The keys defined on this device and on the remote device for the VPN connection. These keys can be different in IKEv2. The key can be 1-127 alphanumeric characters.
- **(IKEv1) Preshared Key**—The key that is defined on both the local and remote device. The key can be 1-127 alphanumeric characters.
- **NAT Exempt**—Whether to exempt the VPN traffic from NAT policies on the local VPN access interface. If you do not want NAT rules to apply to the local network, select the interface that hosts the local network. This option works only if the local network resides behind a single routed interface (not a bridge group member). If the local network is behind more than one routed interface, or one or more bridge group members, you must manually create the NAT exempt rules. For information on manually creating the required rules, see [Exempting Site-to-Site VPN Traffic from NAT, on page 266](#).
- **Diffie-Hellman Group for Perfect Forward Secrecy**—Whether to use Perfect Forward Secrecy (PFS) to generate and use a unique session key for each encrypted exchange. The unique session key protects the exchange from subsequent decryption, even if the entire exchange was recorded and the attacker has obtained the preshared or private keys used by the endpoint devices. To enable Perfect Forward Secrecy, select the Diffie-Hellman key derivation algorithm to use when generating the PFS session key in the Modulus Group list. If you enable both IKEv1 and IKEv2, the options are limited to those supported by IKEv1. For an explanation of the options, see [Deciding Which Diffie-Hellman Modulus Group to Use, on page 252](#).

Step 6 Click **Next**.

Step 7 Review the summary and click **Finish**.

The summary information is copied to the clipboard. You can paste the information in a document and use it to help you configure the remote peer, or to send it to the party responsible for configuring the peer.

After you deploy the configuration, log into the device CLI and use the **show ipsec sa** command to verify that the endpoints establish a security association. See [Verifying Site-to-Site VPN Connections, on page 263](#).

Configuring the Global IKE Policy

Internet Key Exchange (IKE) is a key management protocol that is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and automatically establish IPsec security associations (SAs).

The IKE negotiation comprises two phases. Phase 1 negotiates a security association between two IKE peers, which enables the peers to communicate securely in Phase 2. During Phase 2 negotiation, IKE establishes

SAs for other applications, such as IPsec. Both phases use proposals when they negotiate a connection. An IKE proposal is a set of algorithms that two peers use to secure the negotiation between them. IKE negotiation begins by each peer agreeing on a common (shared) IKE policy. This policy states which security parameters are used to protect subsequent IKE negotiations.

IKE policy objects define the IKE proposals for these negotiations. The objects that you enable are the ones used when the peers negotiate a VPN connection: you cannot specify different IKE policies per connection. The relative priority of each object determines which of these policies are tried first, with the lower number being higher priority. The connection is not established if the negotiation fails to find a policy that both peers can support.

To define the global IKE policy, you select which objects to enable for each IKE version. If the pre-defined objects do not satisfy your requirements, create new policies to enforce your security policy.

The following procedure explains how to configure the global policy through the Objects page. You can also enable, disable, and create policies when editing a VPN connection by clicking **Edit** for the IKE Policy settings.



Note You can enable up to 20 IKE policies.

Procedure

Step 1 Select **Objects**, then select **IKE Policies** from the table of contents.

Policies for IKEv1 and IKEv2 are shown in separate lists.

Step 2 Enable the IKE policies you want to allow for each IKE version.

- a) Select **IKEv1** or **IKEv2** above the object table to show the policies for that version.
- b) Click the **State** toggle to enable the appropriate objects and to disable objects that do not meet your requirements.

If some of your security requirements are not reflected in the existing objects, define new ones to implement your requirements. For details, see the following topics:

- [Configuring IKEv1 Policies, on page 258](#)
- [Configuring IKEv2 Policies, on page 259](#)

- c) Verify that the relative priorities match your requirements.

If you need to change the priority of a policy, edit it. If the policy is a pre-defined system policy, you need to create your own version of the policy to change the priority.

The priority is relative, and not absolute. For example, priority 80 is higher than 160. If 80 is the highest priority object that you enable, that becomes your first-choice policy. If you then enable a policy with priority 25, that becomes your first-choice policy.

- d) If you use both IKE versions, repeat the process for the other version.
-

Configuring IKEv1 Policies

Internet Key Exchange (IKE) version 1 policy objects contain the parameters required for IKEv1 policies when defining VPN connections. IKE is a key management protocol that facilitates the management of IPsec-based communications. It is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and automatically establish IPsec security associations (SAs).

There are several pre-defined IKEv1 policies. If any suit your needs, simply enable them by clicking the **State** toggle. You can also create new policies to implement other combinations of security settings. You cannot edit or delete system-defined objects.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create IKEv1 Policy objects while editing the IKEv1 settings in a VPN connection by clicking the **Create New IKE Policy** link shown in the object list.

Procedure


Step 1 Select **Objects**, then select **IKE Policies** from the table of contents.


Step 2 Select **IKEv1** above the object table to show IKEv1 policies.

Step 3 If any of the system-defined policies meet your requirements, click the **State** toggle to enable them.

Also use the **State** toggle to disable unwanted policies. The relative priority determines which of these policies are tried first, with the lower number being higher priority.

Step 4 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 5 Configure the IKEv1 properties.

- **Priority**—The relative priority of the IKE policy, from 1 to 65,535. The priority determines the order of the IKE policy compared by the two negotiating peers when attempting to find a common security association (SA). If the remote IPsec peer does not support the parameters selected in your highest priority policy, it tries to use the parameters defined in the next lowest priority. The lower the number, the higher the priority.
- **Name**—The name of the object, up to 128 characters.
- **State**—Whether the IKE policy is enabled or disabled. Click the toggle to change the state. Only enabled policies are used during IKE negotiations.
- **Authentication**—The method of authentication to use between the two peers.
 - **Preshared Key**—Use the preshared key that is defined on each device. These keys allow for a secret key to be shared between two peers and to be used by IKE during the authentication phase. If the peer is not configured with the same preshared key, the IKE SA cannot be established.
- **Encryption**—The encryption algorithm used to establish the Phase 1 security association (SA) for protecting Phase 2 negotiations. For an explanation of the options, see [Deciding Which Encryption Algorithm to Use, on page 251](#).

- **Diffie-Hellman Group**—The Diffie-Hellman group to use for deriving a shared secret between the two IPsec peers without transmitting it to each other. A larger modulus provides higher security but requires more processing time. The two peers must have a matching modulus group. For an explanation of the options, see [Deciding Which Diffie-Hellman Modulus Group to Use, on page 252](#).
- **Hash**—The hash algorithm for creating a message digest, which is used to ensure message integrity. For an explanation of the options, see [Deciding Which Hash Algorithms to Use, on page 252](#).
- **Lifetime**—The lifetime of the security association (SA), in seconds, from 120 to 2147483647 or blank. When the lifetime is exceeded, the SA expires and must be renegotiated between the two peers. As a general rule, the shorter the lifetime (up to a point), the more secure your IKE negotiations will be. However, with longer lifetimes, future IPsec security associations can be set up more quickly than with shorter lifetimes. The default is 86400. To specify an unlimited lifetime, enter no value (leave the field blank).

Step 6 Click **OK** to save your changes.

Configuring IKEv2 Policies

Internet Key Exchange (IKE) version 2 policy objects contain the parameters required for IKEv2 policies when defining VPN connections. IKE is a key management protocol that facilitates the management of IPsec-based communications. It is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and automatically establish IPsec security associations (SAs).

There are several pre-defined IKEv2 policies. If any suit your needs, simply enable them by clicking the **State** toggle. You can also create new policies to implement other combinations of security settings. You cannot edit or delete system-defined objects.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create IKEv2 Policy objects while editing the IKEv2 settings in a VPN connection by clicking the **Create New IKE Policy** link shown in the object list.

Procedure


Step 1 Select **Objects**, then select **IKE Policies** from the table of contents.


Step 2 Select **IKEv2** above the object table to show IKEv2 policies.

Step 3 If any of the system-defined policies meet your requirements, click the **State** toggle to enable them.

Also use the **State** toggle to disable unwanted policies. The relative priority determines which of these policies are tried first, with the lower number being higher priority.

Step 4 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 5 Configure the IKEv2 properties.

- **Priority**—The relative priority of the IKE policy, from 1 to 65,535. The priority determines the order of the IKE policy compared by the two negotiating peers when attempting to find a common security association (SA). If the remote IPsec peer does not support the parameters selected in your highest priority policy, it tries to use the parameters defined in the next lowest priority. The lower the number, the higher the priority.
- **Name**—The name of the object, up to 128 characters.
- **State**—Whether the IKE policy is enabled or disabled. Click the toggle to change the state. Only enabled policies are used during IKE negotiations.
- **Encryption**—The encryption algorithm used to establish the Phase 1 security association (SA) for protecting Phase 2 negotiations. Select all algorithms that you want to allow, although you cannot include both mixed-mode (AES-GCM) and normal mode options in the same policy. (Normal mode requires that you select an integrity hash, whereas mixed mode prohibits a separate integrity hash selection.) The system negotiates with the peer, starting from the strongest to the weakest algorithm, until a match is agreed upon. For an explanation of the options, see [Deciding Which Encryption Algorithm to Use, on page 251](#).
- **Diffie-Hellman Group**—The Diffie-Hellman group to use for deriving a shared secret between the two IPsec peers without transmitting it to each other. A larger modulus provides higher security but requires more processing time. The two peers must have a matching modulus group. Select all algorithms that you want to allow. The system negotiates with the peer, starting from the strongest to the weakest group, until a match is agreed upon. For an explanation of the options, see [Deciding Which Diffie-Hellman Modulus Group to Use, on page 252](#).
- **Integrity Hash**—The integrity portion of the hash algorithm for creating a message digest, which is used to ensure message integrity. Select all algorithms that you want to allow. The system negotiates with the peer, starting from the strongest to the weakest algorithm, until a match is agreed upon. The integrity hash is not used with the AES-GCM encryption options. For an explanation of the options, see [Deciding Which Hash Algorithms to Use, on page 252](#).
- **Pseudo Random Function (PRF) Hash**—The pseudo-random function (PRF) portion of the hash algorithm, which is used as the algorithm to derive keying material and hashing operations required for the IKEv2 tunnel encryption. In IKEv1, the Integrity and PRF algorithms are not separated, but in IKEv2, you can specify different algorithms for these elements. Select all algorithms that you want to allow. The system negotiates with the peer, starting from the strongest to the weakest algorithm, until a match is agreed upon. For an explanation of the options, see [Deciding Which Hash Algorithms to Use, on page 252](#).
- **Lifetime**—The lifetime of the security association (SA), in seconds, from 120 to 2147483647 or blank. When the lifetime is exceeded, the SA expires and must be renegotiated between the two peers. As a general rule, the shorter the lifetime (up to a point), the more secure your IKE negotiations will be. However, with longer lifetimes, future IPsec security associations can be set up more quickly than with shorter lifetimes. The default is 86400. To specify an unlimited lifetime, enter no value (leave the field blank).

Step 6 Click **OK** to save your changes.

Configuring IPsec Proposals

IPsec is one of the most secure methods for setting up a VPN. IPsec provides data encryption at the IP packet level, offering a robust security solution that is standards-based. With IPsec, data is transmitted over a public network through tunnels. A tunnel is a secure, logical communication path between two peers. Traffic that enters an IPsec tunnel is secured by a combination of security protocols and algorithms called a transform set. During the IPsec security association (SA) negotiation, peers search for a transform set that is the same at both peers.

There are separate IPsec proposal objects based on the IKE version, IKEv1, or IKEv2:

- When you create an IKEv1 IPsec proposal, you select the mode in which IPsec operates, and define the required encryption and authentication types. You can select single options for the algorithms. If you want to support multiple combinations in a VPN, create and select multiple IKEv1 IPsec Proposal objects.
- When you create an IKEv2 IPsec proposal, you can select all of the encryption and hash algorithms allowed in a VPN. The system orders the settings from the most secure to the least secure and negotiates with the peer until a match is found. This allows you to potentially send a single proposal to convey all the allowed combinations instead of the need to send each allowed combination individually as with IKEv1.

The Encapsulating Security Protocol (ESP) is used for both IKEv1 and IKEv2 IPsec proposals. It provides authentication, encryption, and antireplay services. ESP is IP protocol type 50.



Note We recommend using both encryption and authentication on IPsec tunnels.

The following topics explain how to configure IPsec proposals for each IKE version.

Configuring IPsec Proposals for IKEv1

Use IKEv1 IPsec Proposal objects configure the IPsec proposal used during IKE Phase 2 negotiations. The IPsec proposal defines the combination of security protocols and algorithms that secure traffic in an IPsec tunnel.

There are several pre-defined IKEv1 IPsec proposals. You can also create new proposals to implement other combinations of security settings. You cannot edit or delete system-defined objects.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create IKEv1 IPsec Proposals objects while editing the IKEv1 IPsec settings in a VPN connection by clicking the **Create New IPsec Proposal** link shown in the object list.

Procedure

- Step 1** Select **Objects**, then select **IPsec Proposals** from the table of contents.
- Step 2** Select **IKEv1** above the object table to show IKEv1 IPsec proposals.
- Step 3** Do one of the following:
- To create an object, click the + button.
 - To edit an object, click the edit icon (🔍) for the object.

To delete an unreferenced object, click the trash can icon (🗑️) for the object.

Step 4 Configure the IKEv1 IPsec proposal properties.

- **Name**—The name of the object, up to 128 characters.
- **Mode**—The mode in which the IPsec tunnel operates.
 - **Tunnel** mode encapsulates the entire IP packet. The IPsec header is added between the original IP header and a new IP header. This is the default. Use tunnel mode when the firewall is protecting traffic to and from hosts positioned behind the firewall. Tunnel mode is the normal way regular IPsec is implemented between two firewalls (or other security gateways) that are connected over an untrusted network, such as the Internet.
 - **Transport** mode encapsulates only the upper-layer protocols of an IP packet. The IPsec header is inserted between the IP header and the upper-layer protocol header (such as TCP). Transport mode requires that both the source and destination hosts support IPsec, and can only be used when the destination peer of the tunnel is the final destination of the IP packet. Transport mode is generally used only when protecting a Layer 2 or Layer 3 tunneling protocol such as GRE, L2TP, and DLSW.
- **ESP Encryption**—The Encapsulating Security Protocol (ESP) encryption algorithm for this proposal. For an explanation of the options, see [Deciding Which Encryption Algorithm to Use, on page 251](#).
- **ESP Hash**—The hash or integrity algorithm to use for authentication. For an explanation of the options, see [Deciding Which Hash Algorithms to Use, on page 252](#).

Step 5 Click **OK** to save your changes.

Configuring IPsec Proposals for IKEv2

Use IKEv2 IPsec Proposal objects to configure the IPsec proposal used during IKE Phase 2 negotiations. The IPsec proposal defines the combination of security protocols and algorithms that secure traffic in an IPsec tunnel.

There are several pre-defined IKEv2 IPsec proposals. You can also create new proposals to implement other combinations of security settings. You cannot edit or delete system-defined objects.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create IKEv2 IPsec Proposals objects while editing the IKEv2 IPsec settings in a VPN connection by clicking the **Create New IPsec Proposal** link shown in the object list.


Procedure

Step 1 Select **Objects**, then select **IPsec Proposals** from the table of contents.

Step 2 Select **IKEv2** above the object table to show IKEv2 IPsec proposals.

Step 3 Do one of the following:

- To create an object, click the + button.
- To edit an object, click the edit icon (✎) for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 4 Configure the IKEv2 IPsec proposal properties.

- **Name**—The name of the object, up to 128 characters.
- **Encryption**—The Encapsulating Security Protocol (ESP) encryption algorithm for this proposal. Select all algorithms that you want to allow. The system negotiates with the peer, starting from the strongest to the weakest algorithm, until a match is agreed upon. For an explanation of the options, see [Deciding Which Encryption Algorithm to Use, on page 251](#).
- **Integrity Hash**—The hash or integrity algorithm to use for authentication. Select all algorithms that you want to allow. The system negotiates with the peer, starting from the strongest to the weakest algorithm, until a match is agreed upon. For an explanation of the options, see [Deciding Which Hash Algorithms to Use, on page 252](#).

Note You should choose the null integrity algorithm if you select one of the AES-GCM/GMAC options as the encryption algorithm. These encryption standards do not use the integrity hash even if you select a non-null option.

Step 5 Click **OK** to save your changes.

Verifying Site-to-Site VPN Connections

After you configure a site-to-site VPN connection, and deploy the configuration to the device, verify that the system establishes the security association with the remote device.

If the connection cannot be established, use the **ping interface** *interface_name remote_ip_address* command from the device CLI to ensure there is a path through the VPN interface to the remote device. If there is no connection through the configured interface, you can leave off the **interface** *interface_name* keyword and determine if connectivity is through a different interface. You might have selected the wrong interface for the connection: you must select the interface that faces the remote device, not the interface that faces the protected network.

If there is a network path, check the IKE versions and keys configured and supported by both endpoints, and adjust the VPN connection as needed. Ensure that no access control or NAT rules are blocking the connection.

Procedure

Step 1 Log into the device CLI as explained in [Logging Into the Command Line Interface \(CLI\), on page 5](#).

Step 2 Use the **show ipsec sa** command to verify that the IPsec security association is established.

You should see that the VPN connection is established between your device (the **local addr**) and the remote peer (**current_peer**). The packets (pkts) counts should increase as you send traffic through the connection. The access list should show the local and remote networks for the connection.

For example, the following output shows an IKEv2 connection.

```
> show ipsec sa
interface: site-a-outside
  Crypto map tag: s2sCryptoMap, seq num: 1, local addr: 192.168.2.15
```

```

access-list |s2sAcl|0730e31c-1e5f-11e7-899f-27f6e1030344
extended permit ip 192.168.1.0 255.255.255.0 192.168.3.0 255.255.255.0
local ident (addr/mask/prot/port): (192.168.1.0/255.255.255.0/0/0)
remote ident (addr/mask/prot/port): (192.168.3.0/255.255.255.0/0/0)
current_peer: 192.168.4.6

#pkts encaps: 69, #pkts encrypt: 69, #pkts digest: 69
#pkts decaps: 69, #pkts decrypt: 69, #pkts verify: 69
#pkts compressed: 0, #pkts decompressed: 0
#pkts not compressed: 69, #pkts comp failed: 0, #pkts decomp failed: 0
#pre-frag successes: 0, #pre-frag failures: 0, #fragments created: 0
#PMTUs sent: 0, #PMTUs rcvd: 0, #decapsulated frgs needing reassembly: 0
#TFC rcvd: 0, #TFC sent: 0
#Valid ICMP Errors rcvd: 0, #Invalid ICMP Errors rcvd: 0
#send errors: 0, #recv errors: 0

local crypto endpt.: 192.168.2.15/500, remote crypto endpt.: 192.168.4.6/500
path mtu 1500, ipsec overhead 55(36), media mtu 1500
PMTU time remaining (sec): 0, DF policy: copy-df
ICMP error validation: disabled, TFC packets: disabled
current outbound spi: CD22739C
current inbound spi : 52D2F1E4

inbound esp sas:
spi: 0x52D2F1E4 (1389556196)
SA State: active
transform: esp-aes-gcm-256 esp-null-hmac no compression
in use settings =(L2L, Tunnel, PFS Group 19, IKEv2, )
slot: 0, conn_id: 62738432, crypto-map: s2sCryptoMap
sa timing: remaining key lifetime (kB/sec): (4285434/28730)
IV size: 8 bytes
replay detection support: Y
Anti replay bitmap:
0xFFFFFFFF 0xFFFFFFFF
outbound esp sas:
spi: 0xCD22739C (3441587100)
SA State: active
transform: esp-aes-gcm-256 esp-null-hmac no compression
in use settings =(L2L, Tunnel, PFS Group 19, IKEv2, )
slot: 0, conn_id: 62738432, crypto-map: s2sCryptoMap
sa timing: remaining key lifetime (kB/sec): (4055034/28730)
IV size: 8 bytes
replay detection support: Y
Anti replay bitmap:
0x00000000 0x00000001

```

The following output shows an IKEv1 connection.

```

> show ipsec sa
interface: site-a-outside
Crypto map tag: s2sCryptoMap, seq num: 1, local addr: 192.168.2.15

access-list |s2sAcl|0730e31c-1e5f-11e7-899f-27f6e1030344
extended permit ip 192.168.1.0 255.255.255.0 192.168.3.0 255.255.255.0
local ident (addr/mask/prot/port): (192.168.1.0/255.255.255.0/0/0)
remote ident (addr/mask/prot/port): (192.168.3.0/255.255.255.0/0/0)
current_peer: 192.168.4.6

#pkts encaps: 10, #pkts encrypt: 10, #pkts digest: 10
#pkts decaps: 10, #pkts decrypt: 10, #pkts verify: 10

```



```

#pkts compressed: 0, #pkts decompressed: 0
#pkts not compressed: 10, #pkts comp failed: 0, #pkts decomp failed: 0
#pre-frag successes: 0, #pre-frag failures: 0, #fragments created: 0
#PMTUs sent: 0, #PMTUs rcvd: 0, #decapsulated frgs needing reassembly: 0
#TFC rcvd: 0, #TFC sent: 0
#Valid ICMP Errors rcvd: 0, #Invalid ICMP Errors rcvd: 0
#send errors: 0, #recv errors: 0

local crypto endpt.: 192.168.2.15/0, remote crypto endpt.: 192.168.4.6/0
path mtu 1500, ipsec overhead 74(44), media mtu 1500
PMTU time remaining (sec): 0, DF policy: copy-df
ICMP error validation: disabled, TFC packets: disabled
current outbound spi: 077D72C9
current inbound spi : AC146DEC

inbound esp sas:
spi: 0xAC146DEC (2887020012)
SA State: active
transform: esp-aes-256 esp-sha-hmac no compression
in use settings = {L2L, Tunnel, PFS Group 5, IKEv1, }
slot: 0, conn_id: 143065088, crypto-map: s2sCryptoMap
sa timing: remaining key lifetime (kB/sec): (3914999/28567)
IV size: 16 bytes
replay detection support: Y
Anti replay bitmap:
0x00000000 0x000007FF
outbound esp sas:
spi: 0x077D72C9 (125661897)
SA State: active
transform: esp-aes-256 esp-sha-hmac no compression
in use settings = {L2L, Tunnel, PFS Group 5, IKEv1, }
slot: 0, conn_id: 143065088, crypto-map: s2sCryptoMap
sa timing: remaining key lifetime (kB/sec): (3914999/28567)
IV size: 16 bytes
replay detection support: Y
Anti replay bitmap:
0x00000000 0x00000001

```

Step 3 Use the **show isakmp sa** command to verify the IKE security associations.

You can use the command without the **sa** keyword (or use the **stats** keyword instead) to view IKE statistics.

For example, the following output shows an IKEv2 security association.

```
> show isakmp sa
```

```
There are no IKEv1 SAs
```

```
IKEv2 SAs:
```

```
Session-id:15317, Status:UP-ACTIVE, IKE count:1, CHILD count:1
```

```

Tunnel-id Local Remote Status Role
592216161 192.168.2.15/500 192.168.4.6/500 READY INITIATOR
Encr: AES-GCM, keysize: 256, Hash: N/A, DH Grp:21, Auth sign: PSK, Auth verify: PSK
Life/Active Time: 86400/12 sec
Child sa: local selector 192.168.1.0/0 - 192.168.1.255/65535
remote selector 192.168.3.0/0 - 192.168.3.255/65535
ESP spi in/out: 0x52d2f1e4/0xcd22739c

```

The following output shows an IKEv1 security association.

```

> show isakmp sa

IKEv1 SAs:

    Active SA: 1
    Rekey SA: 0 (A tunnel will report 1 Active and 1 Rekey SA during rekey)
Total IKE SA: 1

1  IKE Peer: 192.168.4.6
   Type      : L2L                Role      : initiator
   Rekey     : no                 State     : MM_ACTIVE

There are no IKEv2 SAs

```

Monitoring Site-to-Site VPN

To monitor and troubleshoot site-to-site VPN connections, log into the device CLI and use the following commands.

- **show ipsec sa** displays the VPN sessions (security associations). You can reset these statistics using the **clear ipsec sa counters** command.
- **show ipsec keyword** displays IPsec operational data and statistics. Enter **show ipsec ?** to see the available keywords.
- **show isakmp** displays ISAKMP operational data and statistics.

Examples for Site-to-Site VPN

The following are examples of configuring site-to-site VPN.

Exempting Site-to-Site VPN Traffic from NAT

When you have a site-to-site VPN connection defined on an interface, and you also have NAT rules for that interface, you can optionally exempt the traffic on the VPN from the NAT rules. You might want to do this if the remote end of the VPN connection can handle your internal addresses.

When you create the VPN connection, you can select the **NAT Exempt** option to create the rules automatically. However, this works only if your local protected network is connected through a single routed interface (not a bridge group member). If instead, the local networks in the connection reside behind two or more routed interfaces, or one or more bridge group members, you need to manually configure the NAT exempt rules.

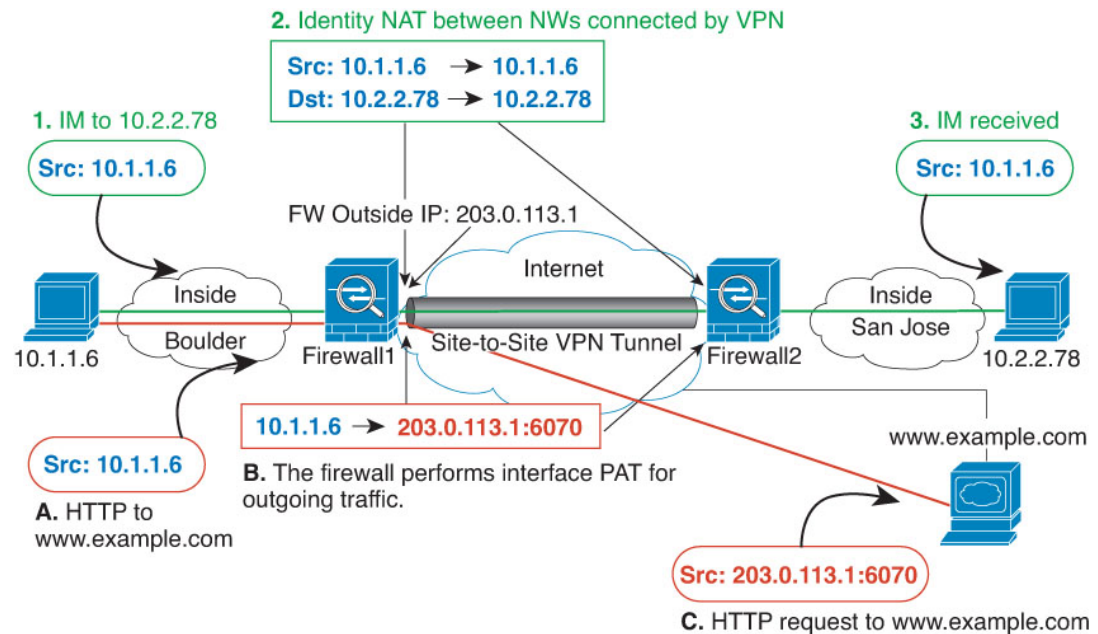
To exempt VPN traffic from NAT rules, you create an identity manual NAT rule for the local traffic when the destination is the remote network. Then, apply NAT to the traffic when the destination is anything else (for example, the Internet). If you have more than one interface for the local network, create rules for each interface. Also, consider the following suggestions:

- If there is more than one local network in the connection, create a network object group to hold the objects that define the networks.

- If you are including both IPv4 and IPv6 networks in the VPN, create separate identity NAT rules for each.

Consider the following example, which shows a site-to-site tunnel connecting the Boulder and San Jose offices. For traffic that you want to go to the Internet (for example from 10.1.1.6 in Boulder to www.example.com), you need a public IP address provided by NAT to access the Internet. The below example uses interface PAT rules. However, for traffic that you want to go over the VPN tunnel (for example from 10.1.1.6 in Boulder to 10.2.2.78 in San Jose), you do not want to perform NAT; you need to exempt that traffic by creating an identity NAT rule. Identity NAT simply translates an address to the same address.

Figure 21: Interface PAT and Identity NAT for Site-to-Site VPN



The following example explains the configuration for Firewall1 (Boulder). The example assumes that the inside interface is a bridge group, so you need to write the rules for each member interface. The process is the same if you have a single or multiple routed inside interfaces.

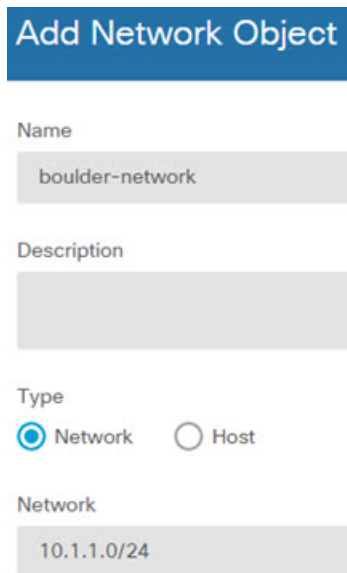


Note This example assumes IPv4 only. If the VPN also includes IPv6 networks, create parallel rules for IPv6. Note that you cannot implement IPv6 interface PAT, so you need to create a host object with a unique IPv6 address to use for PAT.

Procedure

- Step 1** Create the objects to define the various networks.
- Choose **Objects**.
 - Select **Network** from the table of contents and click +.
 - Identify the Boulder inside network.

Name the network object (for example, boulder-network), select **Network**, and enter the network address, 10.1.1.0/24.



Add Network Object

Name
boulder-network

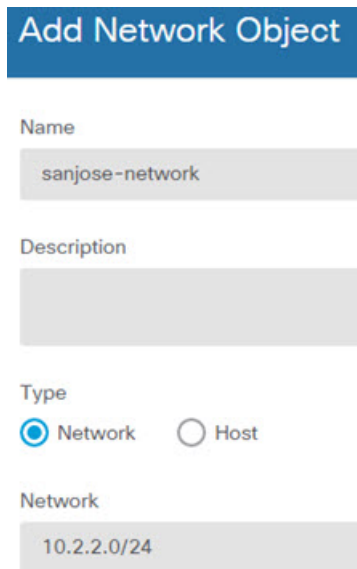
Description

Type
 Network Host

Network
10.1.1.0/24

- d) Click **OK**.
- e) Click + and define the inside San Jose network.

Name the network object (for example, sanjose-network), select **Network**, and enter the network address 10.2.2.0/24.



Add Network Object

Name
sanjose-network

Description

Type
 Network Host

Network
10.2.2.0/24

- f) Click **OK**.

Step 2 Configure manual identity NAT for the Boulder network when going over the VPN to San Jose on Firewall1 (Boulder).

- a) Select **Policies > NAT**.
- b) Click the + button.

c) Configure the following properties:

- **Title** = NAT Exempt 1_2 Boulder San Jose VPN (or another name of your choosing).
- **Create Rule For** = Manual NAT.
- **Placement** = **Above a Specific Rule**, and select the first rule in the Manual NAT Before Auto NAT section. You want to ensure that this rule comes before any general interface PAT rules for the destination interface. Otherwise, the rule might not be applied to the right traffic.
- **Type** = Static.
- **Source Interface** = inside1_2.
- **Destination Interface** = outside.
- **Original Source Address** = boulder-network network object.
- **Translated Source Address** = boulder-network network object.
- **Original Destination Address** = sanjose-network network object.
- **Translated Destination Address** = sanjose-network network object.

Note Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the original and translated destination addresses. Leave all of the port fields blank. This rule configures identity NAT for both source and destination.

- d) On the **Advanced** tab, select **Do not proxy ARP on Destination interface**.
- e) Click **OK**.
- f) Repeat the process to create equivalent rules for each of the other inside interfaces.

Step 3 Configure manual dynamic interface PAT when going to the Internet for the inside Boulder network on Firewall1 (Boulder).

Note There might already be dynamic interface PAT rules for the inside interfaces, covering any IPv4 traffic, as these are created by default during initial configuration. However, the configuration is shown here for completeness. Before completing these steps, check whether a rule already exists that covers the inside interface and network, and skip this step if it does.

- a) Click the + button.
- b) Configure the following properties:
 - **Title** = inside1_2 interface PAT (or another name of your choosing).
 - **Create Rule For** = Manual NAT.
 - **Placement** = **Below a Specific Rule**, and select the rule you created above for this interface in the Manual NAT Before Auto NAT section. Because this rule will apply to any destination address, the rule that uses sanjose-network as the destination must come before this rule, or the sanjose-network rule will never be matched. The default is to place new manual NAT rules at the end of the "NAT Rules Before Auto NAT" section, which is also sufficient.

- **Type** = Dynamic.
- **Source Interface** = inside1_2.
- **Destination Interface** = outside.
- **Original Source Address** = boulder-network network object.
- **Translated Source Address** = **Interface**. This option configures interface PAT using the destination interface.
- **Original Destination Address** = any.
- **Translated Destination Address** = any.

Add NAT Rule ? ×

Title

Create Rule for

Status

Manual NAT rules allow the translation of the source as well as the destination address of a network packet. Destination and port translation are optional. You can place manual NAT rules either before or after Auto NAT rules and insert the rules at a specific location.

Placement

Type

Packet Translation

Advanced Options

ORIGINAL PACKET		TRANSLATED PACKET	
<p>Source Interface <input type="text" value="inside1_2"/></p>	<p>Destination Interface <input type="text" value="outside"/></p>		
<p>Source Address <input type="text" value="boulder-network"/></p>	<p>Source Port <input type="text" value="Any"/></p>	<p>Source Address <input type="text" value="Interface"/></p>	<p>Source Port <input type="text" value="Any"/></p>
<p>Destination Address <input type="text" value="Any"/></p>	<p>Destination Port <input type="text" value="Any"/></p>	<p>Destination Address <input type="text" value="Any"/></p>	<p>Destination Port <input type="text" value="Any"/></p>

- c) Click **OK**.
- d) Repeat the process to create equivalent rules for each of the other inside interfaces.

Step 4 Commit your changes.

- a) Click the **Deploy Changes** icon in the upper right of the web page.



- b) Click the **Deploy Now** button.

Wait for deployment to finish. The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.


Step 5 If you are also managing Firewall2 (San Jose), you can configure similar rules for that device.

- The manual identity NAT rule would be for sanjose-network when the destination is boulder-network. Create new interface objects for the Firewall2 inside and outside networks.
- The manual dynamic interface PAT rule would be for sanjose-network when the destination is "any."

How to Provide Internet Access on the Outside Interface for External Site-to-Site VPN Users (Hair Pinning)

In a site-to-site VPN, you might want users on the remote networks to access the Internet through your device. However, because the remote users are entering your device on the same interface that faces the Internet (the outside interface), you need to bounce Internet traffic right back out of the outside interface. This technique is sometimes called hair pinning.

The following graphic shows an example. There is a site-to-site VPN tunnel configured between 198.51.100.1 (on the main site, Site A) and 203.0.113.1 (the remote site, Site B). All user traffic from the remote site inside network, 192.168.2.0/24, goes through the VPN. Thus, when a user on that network wants to go to a server on the Internet, such as www.example.com, the connection first goes through the VPN, then gets routed back out to the Internet from the 198.51.100.1 interface.



The following procedure explains how to configure this service. You must configure both endpoints of the VPN tunnel.

Procedure

- Step 1** (Site A, main site.) Configure the site-to-site VPN connection to remote Site B.
- Click **Device**, then click **View Configuration** in the Site-to-Site VPN group.
 - Click + to add a new connection.
 - Define the endpoints as follows, and then click **Next**:

- **Connection Profile Name**—Give the connection a meaningful name, for example, Site-A-to-Site-B.
- **Local VPN Access Interface**—Select the outside interface.
- **Local Network**—Keep the default, Any.
- **Remote IP Address**—Enter the IP address of the remote peer’s outside interface. In this example, 203.0.113.1.
- **Remote Network**—Click +, then select the network object that defines the remote peer’s protected network. In this example, 192.168.2.0/24. You can click **Create New Network** to create the object now.

The following graphic shows how the first step should look.

Connection Profile Name

Site-A-to-Site-B

LOCAL SITE	REMOTE SITE
Local VPN Access Interface	Remote IP Address
outside	203.0.113.1
Local Network	Remote Network
+ ANY	+ Site-B-Network

d) Define the privacy configuration, then click **Next**.

- **IKE Policy**—The IKE settings have no impact on hair pinning. Simply select the IKE versions, policies, and proposals that fit your security needs. Make note of the local and remote pre-shared keys you enter: you will need these when configuring the remote peer.
- **NAT Exempt**—Select the inside interface.

Additional Options

NAT Exempt

inside

- **Diffie Helman Group for Perfect Forward Secrecy**—This setting has no impact on hair pinning. Configure it as you see fit.

e) Click **Finish**.

The connection summary is copied to the clipboard. You can paste it into a text file or other document to help you configure the remote peer.

Step 2 (Site A, main site.) Configure the NAT rule to translate all connections going out the outside interface to ports on the outside IP address (interface PAT).

When you complete the initial device configuration, the system creates a NAT rule named `InsideOutsideNatRule`. This rule applies interface PAT to IPv4 traffic from any interface that exits the device through the outside interface. Because the outside interface is included in “Any” source interface, the rule you need already exists, unless you edited it or deleted it.

The following procedure explains how to create the rule you need.

- a) Click **Policies > NAT**.
- b) Do one of the following:
 - To edit the `InsideOutsideNatRule`, mouse over the **Action** column and click the edit icon (🔗).
 - To create a new rule, click +.
- c) Configure a rule with the following properties:
 - **Title**—For a new rule, enter a meaningful name without spaces. For example, `OutsideInterfacePAT`.
 - **Create Rule For**—**Manual NAT**.
 - **Placement**—**Before Auto NAT Rules** (the default).
 - **Type**—**Dynamic**.
 - **Original Packet**—For **Source Address**, select either Any or any-ipv4. For **Source Interface**, ensure that you select Any (which is the default). For all other Original Packet options, keep the default, Any.
 - **Translated Packet**—For **Destination Interface**, select outside. For **Translated Address**, select **Interface**. For all other Translated Packet options, keep the default, Any.

The following graphic shows the simple case where you select Any for the source address.

The screenshot shows the configuration for a Manual NAT rule. Key elements highlighted with red circles include:

- Title:** Create Rule for (dropdown), Manual NAT (dropdown), and Status (toggle switch).
- Placement:** Before Auto NAT Rules (dropdown) and Type: Dynamic (dropdown).
- Packet Translation:**
 - ORIGINAL PACKET:** Source Interface: Any (dropdown), Source Address: Any (dropdown).
 - TRANSLATED PACKET:** Destination Interface: outside (dropdown), Source Address: Interface (dropdown).

d) Click **OK**.

Step 3 (Site A, main site.) Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.



b) Click the **Deploy Now** button.

Wait for deployment to finish. The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.

Step 4 (Site B, remote site.) Log into the remote site's device, and configure the site-to-site VPN connection to Site A.

Use the connection summary obtained from the Site A device configuration to help you configure the Site B side of the connection.

a) Click **Device**, then click **View Configuration** in the Site-to-Site VPN group.

b) Click + to add a new connection.

c) Define the endpoints as follows, and then click **Next**:

- **Connection Profile Name**—Give the connection a meaningful name, for example, Site-B-to-Site-A.
- **Local VPN Access Interface**—Select the outside interface.

- **Local Network**—Click +, then select the network object that defines the local protected network. In this example, 192.168.2.0/24. You can click **Create New Network** to create the object now.
- **Remote IP Address**—Enter the IP address of the main site’s outside interface. In this example, 198.51.100.1.
- **Remote Network**—Keep the default, Any. Ignore the warning; it is not relevant for this use case.

The following graphic shows how the first step should look.

Connection Profile Name

Site-B-to-Site-A

LOCAL SITE	REMOTE SITE
Local VPN Access Interface	Remote IP Address
outside	198.51.100.1
Local Network	Remote Network
+ ProtectedNetwork	i We don't recommend to use "ANY" for this option. + ANY

d) Define the privacy configuration, then click **Next**.

- **IKE Policy**—The IKE settings have no impact on hair pinning. Configure the same or compatible options as those on Site A’s end of the VPN connection. You must configure the pre-shared keys correctly: switch the local and remote keys (for IKEv2) as configured on the Site A device. For IKEv1, there is just one key, which must be the same on both peers.
- **NAT Exempt**—Select the inside interface.

Additional Options

NAT Exempt



inside

- **Diffie Helman Group for Perfect Forward Secrecy**—This setting has no impact on hair pinning. Match the setting used on Site A’s end of the VPN connection.

e) Click **Finish**.

Step 5 (Site B, remote site.) Delete all NAT rules for the protected network so that all traffic leaving the site must go through the VPN tunnel.

Performing NAT on this device is unnecessary because the Site A device will do the address translation. But please examine your specific situation. If you have multiple internal networks and not all of them are participating in this VPN connection, do not delete NAT rules that you need for those networks.

- a) Click **Policies > NAT**.
- b) Do one of the following:
 - To delete rules, mouse over the Action column and click the delete icon ().
 - To edit rules so they no longer apply to the protected network, mouse over the Action column and click the edit icon (.

Step 6 (Site B, remote site.) Commit your changes.

- a) Click the **Deploy Changes** icon in the upper right of the web page.



- b) Click the **Deploy Now** button and wait for deployment to finish.

Wait for deployment to finish. The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.



PART **IV**

System Administration

- [System Settings, on page 281](#)
- [System Management, on page 291](#)



CHAPTER 12

System Settings

The following topics explain how to configure the various system settings that are grouped together on the System Settings page. The settings cover overall system function.

- [Configuring the Management Access List, on page 281](#)
- [Configuring Diagnostic Logging, on page 283](#)
- [Configuring the DHCP Server, on page 284](#)
- [Configuring DNS, on page 286](#)
- [Configuring the Management Interface, on page 286](#)
- [Configuring the Device Hostname, on page 288](#)
- [Configuring Network Time Protocol \(NTP\), on page 288](#)
- [Configuring URL Filtering Preferences, on page 289](#)
- [Configuring Cloud Management \(Cisco Defense Orchestrator\), on page 289](#)

Configuring the Management Access List

By default, you can reach the device's FDM web or CLI interfaces on the management address from any IP address. System access is protected by username/password only. However, you can configure an access list to allow connections from specific IP addresses or subnets only to provide another level of protection.

You can also open data interfaces to allow the FDM or SSH connections to the CLI. You can then manage the device without using the management address. For example, you could allow management access to the outside interface, so that you can configure the device remotely. The username/password protects against unwanted connections. By default, HTTPS management access to data interfaces is enabled on the inside interface but it is disabled on the outside interface. For any device model that has a default “inside” bridge group, this means that you can make the FDM connections through any data interface within the bridge group to the bridge group IP address (default is 192.168.1.1). You can open a management connection only on the interface through which you enter the device.



Caution

If you constrain access to specific addresses, you can easily lock yourself out of the system. If you delete access for the IP address that you are currently using, and there is no entry for “any” address, you will lose access to the system when you deploy the policy. Be very careful if you decide to configure the access list.

Procedure

Step 1 Click **Device**, then click the **System Settings > Management Access List** link.

If you are already on the System Settings page, simply click **Management Access List** in the table of contents.

Step 2 To create rules for the management address:

a) Select the **Management Interface** tab.

The list of rules defines which addresses are allowed access to the indicated port: 443 for the FDM (the HTTPS web interface), 22 for the SSH CLI.

The rules are not an ordered list. If an IP address matches any rule for the requested port, the user is allowed to attempt logging into the device.

Note To delete a rule, click the trash can icon (🗑️) for the rule. If you delete all of the rules for a protocol, no one can access the device on that interface using the protocol.

b) Click + and fill in the following options:

- **Protocol**—Select whether the rule is for HTTPS (port 443) or SSH (port 22).
- **IP Address**—Select the network object that defines the IPv4 or IPv6 network or host that should be able to access the system. To specify "any" address, select **any-ipv4** (0.0.0.0/0) and **any-ipv6** (::/0).

c) Click **OK**.

Step 3 To create rules for data interfaces:

a) Select the **Data Interfaces** tab.

The list of rules defines which addresses are allowed access to the indicated port on the interface: 443 for the FDM (the HTTPS web interface), 22 for the SSH CLI.

The rules are not an ordered list. If an IP address matches any rule for the requested port, the user is allowed to attempt logging into the device.

Note To delete a rule, click the trash can icon (🗑️) for the rule. If you delete all of the rules for a protocol, no one can access the device on that interface using the protocol.

b) Click + and fill in the following options:

- **Interface**—Select the interface on which you want to allow management access.
- **Protocols**—Select whether the rule is for HTTPS (port 443), SSH (port 22), or both.
- **Allowed Networks**—Select the network objects that define the IPv4 or IPv6 network or host that should be able to access the system. To specify "any" address, select **any-ipv4** (0.0.0.0/0) and **any-ipv6** (::/0).

c) Click **OK**.

Configuring Diagnostic Logging

Diagnostic logging provides syslog messages for events that are not related to connections. You configure connection logging within individual access control rules. The following procedure explains how to configure the logging of diagnostic messages.

Procedure

-
- Step 1** Click **Device**, then click the **System Settings > Logging Settings** link.
If you are already on the System Settings page, simply click **Logging Settings** in the table of contents
- Step 2** Click **Diagnostic Log Settings > On**.
Even if you configure the remaining fields on this page, diagnostic log messages are not generated unless you turn on this setting.
- Step 3** Turn the slider to **On** for each of the locations where you want to see diagnostic log messages, and select a minimum severity level.
You can log messages to the following locations:
- **Console**—These messages appear when you log into the CLI on the Console port. You can also see these logs in an SSH session to other interfaces (including the management address) by using the **show console-output** command. In addition, you can see these messages in real time in the diagnostic CLI, enter **system support diagnostic-cli** from the main CLI.
 - **Syslog**—These messages are sent to the external syslog servers that you specify. Click **+**, select the syslog server objects, and click **OK** in the popup dialog box. If the object for a server does not already exist, click **Add Syslog Server** to create it.
- Step 4** Click **Save**.
-

Severity Levels

The following table lists the syslog message severity levels.

Table 5: Syslog Message Severity Levels

Level Number	Severity Level	Description
0	emergencies	System is unusable.
1	alert	Immediate action is needed.
2	critical	Critical conditions.
3	error	Error conditions.
4	warning	Warning conditions.

Level Number	Severity Level	Description
5	notification	Normal but significant conditions.
6	informational	Informational messages only.
7	debugging	Debugging messages only. Log at this level only temporarily, when debugging issues. This log level can potentially generate so many messages that system performance can be affected.



Note ASA and FTD do not generate syslog messages with a severity level of zero (emergencies).

Configuring the DHCP Server

A DHCP server provides network configuration parameters, such as IP addresses, to DHCP clients. You can configure a DHCP server on an interface to provide configuration parameters to DHCP clients on the attached network.

An IPv4 DHCP client uses a broadcast rather than a multicast address to reach the server. The DHCP client listens for messages on UDP port 68; the DHCP server listens for messages on UDP port 67. The DHCP server does not support BOOTP requests.

DHCP clients must be on the same network as the interface on which the server is enabled. That is, there cannot be an intervening router between the server and client, although there can be a switch.



Note Do not configure a DHCP server on a network that already has a DHCP server operating on it. The two servers will conflict and results will be unpredictable.

Procedure

Step 1 Click **Device**, then click the **System Settings > DHCP Server** link.

If you are already on the System Settings page, simply click **DHCP Server** in the table of contents.

The page has two tabs. Initially, the **Configuration** tab shows the global parameters.

The **DHCP Servers** tab shows the interfaces on which you have configured DHCP server, whether the server is enabled, and the address pool for the server.


Step 2 On the **Configuration** tab, configure auto-configuration and global settings.

DHCP auto configuration enables the DHCP Server to provide DHCP clients with DNS server, domain name, and WINS server information obtained from a DHCP client that is running on the specified interface. Typically, you would use auto-configuration if you are obtaining an address using DHCP on the outside interface, but

you could choose any interface that obtains its address through DHCP. If you cannot use auto-configuration, you can manually define the required options.

- a) Click **Enable Auto Configuration** > **On** (the slider should be on the right) if you want to use auto-configuration, and then select the interface that is obtaining its address through DHCP in **From Interface**.
- b) If you do not enable auto-configuration, or if you want to override any of the automatically configured settings, configure the following global options. These settings will be sent to DHCP clients on all interfaces that host DHCP server.
 - **Primary WINS IP Address, Secondary WINS IP Address**—The addresses of the Windows Internet Name Service (WINS) servers clients should use for NetBIOS name resolution.
 - **Primary DNS IP Address, Secondary DNS IP Address**—The addresses of the Domain Name System (DNS) servers clients should use for domain name resolution. Click **Use OpenDNS** if you want to configure the OpenDNS public DNS servers. Clicking the button loads the appropriate IP addresses into the fields.
- c) Click **Save**.

Step 3 Click the **DHCP Servers** tab and configure the servers.

- a) Do one of the following:
 - To configure DHCP server for an interface that is not already listed, click +.
 - To edit an existing DHCP server, click the edit icon () for the server.

To delete a server, click the trash can icon () for the server.

- b) Configure the server properties:
 - **Enable DHCP Server**—Whether to enable the server. You can configure a server but keep it disabled until you are ready to use it.
 - **Interface**—Select the interface on which you will provide DHCP addresses to clients. The interface must have a static IP address; you cannot be using DHCP to obtain the interface address if you want to run a DHCP server on the interface. For bridge groups, you configure the DHCP server on the Bridge Virtual Interface (BVI), not the member interfaces, and the server operates on all member interfaces.

You cannot configure DHCP server on the Diagnostic interface; configure it on the Management interface instead, on the **Device** > **System Settings** > **Management Interface** page.
 - **Address Pool**—The range of IP addresses from lowest to highest that the server is allowed to provide to clients that request an address. Specify the start and end address for the pool, separated by a hyphen. For example, 10.100.10.12-10.100.10.250.

The range of IP addresses must be on the same subnet as the selected interface and cannot include: the IP address of the interface itself, the broadcast address, or the subnet network address.

The size of the address pool is limited to 256 addresses per pool on the FTD device. If the address pool range is larger than 253 addresses, the netmask of the FTD interface cannot be a Class C address (for example, 255.255.255.0) and needs to be something larger, for example, 255.255.254.0.

- c) Click **OK**.
-

Configuring DNS

The Domain Name System (DNS) servers are used to resolve hostnames to IP addresses. These servers are used by the management interface. You configure DNS servers during initial system setup, but you can change them using the following procedure.

You can also change the DNS configuration in the CLI using the **configure network dns servers** and **configure network dns searchdomains** commands.

If you have problems with DNS resolution, see [Troubleshooting DNS for the Management Interface, on page 305](#).

Procedure

- Step 1** Click **Device**, then click the **System Settings > DNS Server** link.
- If you are already on the System Settings page, simply click **DNS Server** in the table of contents.
- Step 2** In **Primary, Secondary, Tertiary DNS IP address**, enter the IP addresses of up to three DNS servers in order of preference.
- The primary DNS server is used unless it cannot be contacted, in which case the secondary is tried, and finally the tertiary.
- Click **Use OpenDNS** if you want to configure the OpenDNS public DNS servers. Clicking the button loads the appropriate IP addresses into the fields.
- Step 3** In **Domain Search Name**, enter the domain name for your network, e.g. example.com.
- This domain is added to hostnames that are not fully-qualified, for example, serverA instead of serverA.example.com.
- Step 4** Click **Save**.
-

Configuring the Management Interface

The Management interface is a virtual interface attached to the physical Management port. Note that the physical interface also includes the Diagnostic virtual interface, which you can configure on the **Interfaces** page with other physical interfaces. See [Management/Diagnostic Interface, on page 100](#) for more information about the Diagnostic interface.

The management interface has two uses:

- You can open web and SSH connections to the IP address and configure the device through the interface.
- The system obtains smart licensing and database updates through this IP address.

If you use the CLI setup wizard, you configure the management address and gateway for the device during initial system configuration. If you use the FDM setup wizard, the management address and gateway remain the defaults.

If necessary, you can change these addresses through the FDM. You can also change the management address and gateway in the CLI using the **configure network ipv4 manual** and **configure network ipv6 manual** commands.

You can define static addresses, or obtain an address through DHCP if another device on the management network is acting as a DHCP server. By default, the management address is static, and a DHCP server runs on the port. Thus, you can plug a device directly into the management port and get a DHCP address for your workstation. This makes it easy to connect to and configure the device.



Caution If you change the address to which you are currently connected, you will lose access to the FDM (or the CLI) when you save the changes, as they are applied immediately. You will need to reconnect to the device. Ensure that the new address is valid and available on the management network.

Procedure

- Step 1** Click **Device**, then click the **System Settings > Management Interface** link.
- If you are already on the **System Settings** page, click **Management Interface** in the table of contents
- Step 2** Choose how you want to define the management gateway.
- The gateway determines how the system can reach the internet to obtain smart licenses, database updates (such as VDB, rule, Geolocation, URL), and to reach the management DNS and NTP servers. Choose from these options:
- (Static IP only) **Use the Data Interfaces as the Gateway**—Select this option if you do not have a separate management network connected to the Management interface. Traffic is routed to the internet based on the routing table, typically going through the outside interface.
 - **Use Unique Gateways for the Management Interface**—Specify unique gateways (below) for IPv4 and IPv6 if you have a separate management network connected to the Management interface. For DHCP IP addressing, the gateway is provided by the DHCP server.
- Step 3** Configure the management address, subnet mask or IPv6 prefix, and gateway (if necessary) for IPv4, IPv6, or both.
- You must configure at least one set of properties. Leave one set blank to disable that addressing method.
- Select **Type > DHCP** to obtain the address and gateway through DHCP or IPv6 auto configuration. However, you cannot use DHCP if you are using the data interfaces as the gateway. In this case, you must use a static address.
- Step 4** (Optional.) If you configure a static IPv4 address, configure a DHCP server on the interface.
- If you configure a DHCP server on the management interface, clients on the management network can obtain their address from the DHCP pool.
- a) Click **Enable DHCP Server > On**.
 - b) Enter the **Address Pool** for the server.

The address pool is the range of IP addresses from lowest to highest that the server is allowed to provide to clients that request an address. The range of IP addresses must be on the same subnet as the management address and cannot include: the IP address of the interface itself, the broadcast address, or the subnet network address. Specify the start and end address for the pool, separated by a hyphen. For example, 192.168.45.46-192.168.45.254.

Step 5 Click **Save**, read the warning, and click **OK**.

Configuring the Device Hostname

You can change the device hostname.

You can also change the hostname in the CLI using the **configure network hostname** command.



Caution If you change the hostname when connected to the system using the hostname, you will lose access to the FDM when you save the changes, as they are applied immediately. You will need to reconnect to the device.

Procedure

- Step 1** Click **Device**, then click the **System Settings > Hostname** link.
If you are already on the System Settings page, simply click **Hostname** in the table of contents
- Step 2** Enter a new hostname.
- Step 3** Click **Save**.
-

Configuring Network Time Protocol (NTP)

You must configure Network Time Protocol (NTP) servers to define the time on the system. You configure NTP servers during initial system setup, but you can change them using the following procedure. If you have problems with the NTP connection, see [Troubleshooting NTP, on page 303](#).

The FTD device supports NTPv4.

Procedure

- Step 1** Click **Device**, then click the **System Settings > NTP** link.
If you are already on the System Settings page, simply click **NTP** in the table of contents
- Step 2** In **NTP Time Server**, select whether you want to use your own or Cisco's time servers.

- **Default NTP Time Server**—If you select this option, the server list shows the server names that are used for NTP.
- **Manually Input**—If you select this option, enter the fully qualified domain name or IPv4 or IPv6 address of the NTP server you want to use. For example, ntp1.example.com or 10.100.10.10. You can add up to 3 NTP servers.

Step 3 Click **Save**.

Configuring URL Filtering Preferences

The system obtains the URL category and reputation database from Cisco Collective Security Intelligence (CSI). These preferences control database updates and how the system handles URLs with unknown category or reputation. You must enable the URL filtering license to set these preferences.

Procedure

Step 1 Click **Device**, then click the **System Settings > URL Filtering Preferences** link.

If you are already on the System Settings page, simply click **URL Filtering Preferences** in the table of contents

Step 2 Configure the following options:

- **Enable Automatic Updates**—Allows the system to automatically check for and download updated URL data, which includes category and reputation information. The system checks for updates every 30 minutes, although the data is typically updated once per day. The default is to enable updates. If you deselect this option, and you are using category and reputation filtering, periodically enable it to get new URL data.
- **Query Cisco CSI for Unknown URLs**—Whether to check with Cisco CSI for updated information for URLs that do not have category and reputation data in the local URL filtering database. If the lookup returns this information within a reasonable time limit, it is used when selecting access rules based on URL conditions. Otherwise, the URL matches the Uncategorized category. Selecting this option is important for lower-end systems, which install a smaller URL database due to memory limitations.

Step 3 Click **Save**.

Configuring Cloud Management (Cisco Defense Orchestrator)

You can manage the device using the Cisco Defense Orchestrator (CDO) cloud-based portal.

Using CDO, you can approach device management using the following techniques:

- **Initial configuration download**—In this approach, you download the initial device configuration from CDO, but thereafter you configure the device locally using FDM.



Note After configuring the device using FDM, if you decide you want to instead manage the device through the cloud, ensure that you duplicate your local changes in the cloud-based configuration.

- Remote configuration management through the cloud—In this approach, you use CDO to create and update the device configuration. When using this approach, do not make local changes to the configuration, because on each cloud deployment, the configuration defined in the cloud replaces the local configuration on the device. If you make a local change, be sure to repeat the configuration in the cloud-based configuration if you want to preserve the change.

For more information about how cloud management works, refer to the CDO portal (<http://www.cisco.com/go/cdo>) or ask the reseller or partner with whom you are working.

Before you begin

Obtain a registration key for CDO.

If you have already registered the device with Cisco Smart Software Manager (CSSM), we strongly recommend that you first unregister the device from the Smart Licensing page. You can re-register after you enable CDO using a token.

Also, ensure that the device has a route to the Internet.

Procedure

Step 1 Click **Device**, then click the **System Settings > Cloud Management** link.

If you are already on the System Settings page, simply click **Cloud Management** in the table of contents.

Step 2 Click **Get Started**.

Step 3 Paste the key in **Registration Key** and click **Connect**.

A registration request is sent to the cloud portal. If the key is valid, and there is a route to the Internet, the device should be successfully registered with the portal. You can then start using the portal to manage the device.

If you decide you no longer want to use cloud management, you can select **Unregister** from the gear drop-down list.



CHAPTER 13

System Management

The following topics explain how to perform system management tasks such as updating system databases and backing up and restoring the system.

- [Installing Software Updates, on page 291](#)
- [Backing Up and Restoring the System, on page 295](#)
- [Rebooting the System, on page 299](#)
- [Troubleshooting the System, on page 300](#)
- [Uncommon Management Tasks, on page 310](#)

Installing Software Updates

You can install updates to the system databases and to the system software. The following topics explain how to install these updates.

Updating System Databases

The system uses several databases to provide advanced services. Cisco provides updates to these databases so that your security policies use the latest information available.

Overview of System Database Updates

FTD uses the following databases to provide advanced services.

Intrusion rules

As new vulnerabilities become known, the Cisco Talos Intelligence Group (Talos) releases intrusion rule updates that you can import. These updates affect intrusion rules, preprocessor rules, and the policies that use the rules.

Intrusion rule updates provide new and updated intrusion rules and preprocessor rules, modified states for existing rules, and modified default intrusion policy settings. Rule updates may also delete rules, provide new rule categories and default variables, and modify default variable values.

For changes made by an intrusion rule update to take effect, you must redeploy the configuration.

Intrusion rule updates may be large, so import rules during periods of low network use. On slow networks, an update attempt might fail, and you will need to retry.

Geolocation database (GeoDB)

The Cisco Geolocation Database (GeoDB) is a database of geographical data (such as country, city, coordinates) associated with routable IP addresses.

GeoDB updates provide updated information on physical locations that your system can associate with detected routable IP addresses. You can use geolocation data as a condition in access control rules.

The time needed to update the GeoDB depends on your appliance; the installation usually takes 30 to 40 minutes. Although a GeoDB update does not interrupt any other system functions (including the ongoing collection of geolocation information), the update does consume system resources while it completes. Consider this when planning your updates.

Vulnerability database (VDB)

The Cisco Vulnerability Database (VDB) is a database of known vulnerabilities to which hosts may be susceptible, as well as fingerprints for operating systems, clients, and applications. The firewall system correlates the fingerprints with the vulnerabilities to help you determine whether a particular host increases your risk of network compromise. The Cisco Talos Intelligence Group (Talos) issues periodic updates to the VDB.

The time it takes to update vulnerability mappings depends on the number of hosts in your network map. You may want to schedule the update during low system usage times to minimize the impact of any system downtime. As a rule of thumb, divide the number of hosts on your network by 1000 to determine the approximate number of minutes to perform the update.

After you update the VDB, you must redeploy configurations before updated application detectors and operating system fingerprints can take effect.

URL Category/Reputation Database

The system obtains the URL category and reputation database from Cisco Collective Security Intelligence (CSI). If you configure URL filtering access control rules that filter on category and reputation, requested URLs are matched against the database. You can configure database updates and some other URL filtering preferences on **System Settings > URL Filtering Preferences**. You cannot manage URL category/reputation database updates the same way you manage updates for the other system databases.

Updating System Databases

You can manually retrieve and apply system database updates at your convenience. Updates are retrieved from the Cisco support site. Thus, there must be a path to the internet from the system's management address.



Note In May 2022 we split the GeoDB into two packages: a country code package that maps IP addresses to countries/continents, and an IP package that contains additional contextual data associated with routable IP addresses. The FDM does not and has never used the information in the IP package. This split saves significant disk space in locally managed FTD deployments. If you are getting the GeoDB from Cisco yourself, make sure you get the country code package, which has the same file name as the old all-in-one package: `Cisco_GEODB_Update-date-build`.

You can also set up a regular schedule to retrieve and apply database updates. Because these updates can be large, schedule them for times of low network activity.



Note While a database update is in progress, you might find that the user interface is sluggish to respond to your actions.

Before you begin

To avoid any potential impact to pending changes, deploy the configuration to the device before manually updating these databases.

Please be aware that VDB and URL category updates can remove applications or categories. You need to update any access control or SSL decryption rules that use these deprecated items before you can deploy changes.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Updates summary.

This opens the Updates page. Information on the page shows the current version for each database and the last date and time each database was updated.

Step 2 To manually update a database, click **Update Now** in the section for that database.

After downloading and applying the update, the system automatically re-deploys policies to the device so that the system can use the updated information.

Step 3 (Optional) To set up a regular database update schedule:

a) Click the **Configure** link in the section for the desired database. If there is already a schedule, click **Edit**.

The update schedules for the databases are separate. You must define the schedules separately.

b) Set the update start time:

- The frequency of the update (Daily, Weekly, or Monthly).
- For weekly or monthly, the days of the week or month you want the update to occur.
- The time you want the update to start. The time you specify is adjusted for Daylight Savings Time, so it will move an hour forward or backward whenever the time is adjusted in your area. You must edit the schedule at the time change if you want to keep this exact time throughout the year.

c) Click **Save**.

Note If you want to remove a recurring schedule, click the **Edit** link to open the scheduling dialog box, then click the **Remove** button.

Upgrading FTD Software

You can install the FTD software upgrades as they become available. The following procedure assumes that your system is already running the FTD version 6.2.0 or higher and that it is operating normally.

Upgrades can be major (A.x), maintenance release (A.x.y), or patch (A.x.y.z). We also may provide hotfixes, which are minor updates that address particular, urgent issues. A hotfix might not require a reboot, while the other upgrade types do require a reboot. The system automatically reboots after installation if a reboot is required. Installing any update can disrupt traffic, so do the installation in off hours.

If you also need to upgrade the FXOS software on the chassis, install the FXOS upgrade before following this procedure.

You cannot reimage a device, or migrate from ASA software to FTD software, using this procedure.



Note Before installing an update, make sure that you deploy any pending changes. You should also run a backup and download the backup copy. Note that all upgrades except hot fixes will delete all backup files retained on the system.

Before you begin

Check the task list and verify there are no tasks running. Please wait until all tasks, such as database updates, are completed before you install an upgrade. Also, check for any scheduled tasks. No scheduled tasks should overlap with the upgrade task.

Prior to performing an update, ensure that no deprecated applications are present in application filters, access rules, or SSL decryption rules. These applications have "(Deprecated)" following the application name. While it is not possible to add deprecated applications to these objects, a subsequent VDB update can cause previously valid applications to become deprecated. If this happens, the upgrade will fail, leaving the device in an unusable state.

Download upgrade files from the Cisco Support & Download site: <https://www.cisco.com/go/ftd-software>.

- Ensure that you obtain the appropriate upgrade file, whose file type is .sh. Do not download the system software package or the boot image.
- Do not rename the upgrade file. The system considers renamed files to be invalid.
- You cannot downgrade or uninstall a patch.
- Verify that you are running the required baseline image for the upgrade. For compatibility information, see the [Cisco Secure Firewall Threat Defense Compatibility Guide](#).
- Read the [Cisco Firepower Release Notes](#) for the new version.

Procedure

Step 1 Select **Device**, then click **View Configuration** in the Updates summary.

The **System Upgrade** section shows the currently running software version and any update that you have already uploaded.

Step 2 Upload the upgrade file.

- If you have not yet uploaded an upgrade file, click **Browse** and select the file.
- If there is already an uploaded file, but you want to upload a different one, click the **Upload Another File** link. You can upload one file only. If you upload a new file, it replaces the old file.

- To remove the file, click the delete icon (🗑).

Step 3 Click **Install** to start the installation process.

Information next to the icon indicates whether the device will reboot during installation. You are automatically logged out of the system. Installation might take 30 minutes or more.

Wait before logging into the system again. The Device Summary, or System monitoring dashboard, should show the new version.

Note Do not simply refresh the browser window. Instead, delete any path from the URL, and reconnect to the home page. This ensures that cached information gets refreshed with the latest code.

Step 4 (Optional.) Update the system databases.

If you do not have automatic update jobs configured for Geolocation, Rule, and Vulnerability (VDB) databases, this is a good time to update them.

Reimaging the Device

Reimaging a device involves wiping out the device configuration and installing a fresh software image. The intention of reimaging is to have a clean installation with a factory default configuration.

You would reimage the device in these circumstances:

- You want to convert the system from ASA Software to FTD Software. You cannot upgrade a device running an ASA image to one running a FTD image.
- The device is running a pre-6.1.0 image, and you want to upgrade to 6.1 or a later image and configure the device using the FDM. You cannot use the FMC to upgrade a pre-6.1 device and then switch to local management.
- The device is not functioning correctly and all attempts at fixing the configuration have failed.

For information on how to reimage a device, see *Reimage the Cisco ASA or Threat Defense Device* or the *Threat Defense Quick Start* guide for your device model. These guides are available at <http://www.cisco.com/c/en/us/support/security/firepower-ngfw/products-installation-guides-list.html>.

Backing Up and Restoring the System

You can back up the system configuration so that you can restore the device if the configuration becomes corrupted due to subsequent miss-configuration or physical mishap.

You can restore a backup onto a replacement device only if the two devices are the same model and are running the same version of the software (including the build number, not just the same point release). Do not use the backup and restore process to copy configurations between appliances. A backup file contains information that uniquely identifies an appliance, so that it cannot be shared in this manner.



Note The backup does not include the management IP address configuration. Thus, when you recover a backup file, the management address is not replaced from the backup copy. This ensures that any changes you made to the address are preserved, and also makes it possible to restore the configuration on a different device on a different network segment.

Backups include the configuration only, and not the system software. If you need to completely reimage the device, you need to reinstall the software, then you can upload a backup and recover the configuration.

The configuration database is locked during backup. You cannot make configuration changes during a backup, although you can view policies, dashboards, and so forth. During a restore, the system is completely unavailable.

The table on the Backup and Restore page lists all existing backup copies that are available on the system, including the file name of the backup, the date and time it was created, and the file size. The type of backup (manual, scheduled, or recurring) is based on how you directed the system to create that backup copy.



Tip Backup copies are created on the system itself. You must manually download backup copies and store them on secure servers to ensure that you have the backup copies you need for disaster recovery. The system maintains up to 3 backup copies on the device. New backups replace the oldest backup.

The following topics explain how to manage backup and restore operations.

Backing Up the System Immediately

You can start a backup whenever you want.

Procedure

- Step 1** Click **Device**, then click **View Configuration** in the Backup and Restore summary.
- This opens the Backup and Restore page. The table lists all existing backup copies that are available on the system.
- Step 2** Click **Manual Backup > Back Up Now**.
- Step 3** Enter a name for the backup and optionally a description.
- If you decide you want to perform the backup at a future time rather than immediately, you can click **Schedule** instead.
- Step 4** Click **Back Up Now**.
- The system starts the backup process. When the backup is complete, the backup file will appear in the table. You can then download the backup copy to your system and store it elsewhere, if desired.
- You can leave the Backup and Restore page after initiating the backup. However, the system will likely be sluggish, and you should consider pausing your work to allow the backup to complete.

In addition, the system will acquire locks on the configuration database during parts or all of the backup, which can prevent you from making changes for the duration of the backup process.

Backing Up the System at a Scheduled Time

You can set up a scheduled backup to back up the system at a specific future date and time. A scheduled backup is a one-time occurrence. If you want to create a backup schedule to regularly create backups, configure a recurring backup instead of a scheduled backup.



Note If you want to delete the schedule for a future backup, edit the schedule and click **Remove**.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Backup and Restore summary.

Step 2 Click **Scheduled Backup** > **Schedule a Backup**.

If you already have a scheduled backup, click **Scheduled Backup** > **Edit**.

Step 3 Enter a name for the backup and optionally a description.

Step 4 Select the date and time for the backup.

Step 5 Click **Schedule**.

When the selected date and time arrives, the system takes a backup. When completed, the backup copy is listed in the table of backups.

Setting Up a Recurring Backup Schedule

You can set up a recurring backup to back up the system on a regular schedule. For example, you could take a backup every Friday at midnight. A recurring backup schedule helps ensure that you always have a set of recent backups.



Note If you want to delete a recurring schedule, edit the schedule and click **Remove**.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Backup and Restore summary.

Step 2 Click **Recurring Backup** > **Configure**.

If you already have a recurring backup configured, click **Recurring Backup** > **Edit**.

Step 3 Enter a name for the backup and optionally a description.

Step 4 Select the **Frequency** and the related schedule:

- **Daily**—Select the time of day. A backup is taken every day at the scheduled time.
- **Weekly**—Select the days of the week and the time of day. A backup is taken on each day you select at the scheduled time. For example, you could schedule backups for every Monday, Wednesday, and Friday at 23:00 hours (11 PM).
- **Monthly**—Select the days of the month and the time of day. A backup is taken on each day you select at the scheduled time. For example, you could schedule backups for the first (1), fifteenth (15), and twenty-eighth (28) days of the month at 23:00 hours (11 PM).

The time you specify is adjusted for Daylight Savings Time, so it will move an hour forward or backward whenever the time is adjusted in your area. You must edit the schedule at the time change if you want to keep this exact time throughout the year.

Step 5 Click **Save**.

When the selected dates and times arrive, the system takes a backup. When completed, the backup copy is listed in the table of backups.

The recurring schedule continues to take backups until you change or remove it.

Restoring a Backup

You can restore backups as needed so long as the device is running the same software version (including build number) as it was running when you took the backup. You can restore a backup onto a replacement device only if the two devices are the same model and are running the same version of the software (including build number).

If the backup copy you want to restore is not already on the device, you must upload the backup first before restoring it.

During a restore, the system is completely unavailable.




Note The backup does not include the management IP address configuration. Thus, when you recover a backup file, the management address is not replaced from the backup copy. This ensures that any changes you made to the address are preserved, and also makes it possible to restore the configuration on a different device on a different network segment.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Backup and Restore summary.

This opens the Backup and Restore page. The table lists all existing backup copies that are available on the system.

Step 2 If the backup copy that you want to restore is not in the list of available backups, click **Upload > Browse** and upload the backup copy.

Step 3 Click the restore icon () for the file.

You are asked to confirm the restore. By default, the backup copy will be deleted after the restore, but you can select **Do not remove the backup after restoring** to keep it before proceeding with the restore.


The system will reboot after restore completes.

Note After the system reboots, it automatically checks for Vulnerability Database (VDB), Geolocation, and Rules database updates, and downloads them if needed. Because these updates can be large, the initial attempt might fail. Please check the task list, and if a download failed, manually download an update as described in [Updating System Databases, on page 292](#). The system also redeploys policies. Any subsequent deployment will fail until the update is successful.

Managing Backup Files



As you create new backups, the backup files are listed on the Backup and Restore page. Backup copies are not retained indefinitely: as disk space usage on the device reaches the maximum threshold, older backup copies are deleted to make room for newer ones. In addition, when you install any upgrade other than a hot fix, all backup files are deleted. Thus, you should regularly manage the backup files to ensure that you have the specific backup copies you most want to keep.

You can do the following to manage your backup copies:

- Download files to secure storage—To download a backup file to your workstation, click the download icon () for the file. You can then move the file to your secure file storage.
- Upload a backup file to the system—If you want to restore a backup copy that is no longer available on the device, click **Upload > Browse File** and upload it from your workstation. You can then restore it.



Note Uploaded files may be renamed to match the original filename. Also, if there are more than 10 backup copies already on the system, the oldest one will be deleted to make room for the uploaded file. You cannot upload files that were created by an older software version.

- Restore a backup—To restore a backup copy, click the restore icon () for the file. The system is unavailable during restore, and will reboot after restore completes. You should deploy the configuration after the system is up and running.
- Delete a backup file—If you no longer want a particular backup, click the delete icon () for the file. You are asked to confirm the deletion. Once deleted, you cannot recover the backup file.

Rebooting the System

If you believe the system is not performing correctly and other efforts to resolve the problem have failed, you can reboot the device. You must reboot the device through the CLI; you cannot reboot the device through the FDM.

Procedure

Step 1 Use an SSH client to open a connection to the management IP address and log into the device CLI with a username that has configuration CLI access. For example, the **admin** username.

Step 2 Enter the **reboot** command.

Example:

```
> reboot
```

Troubleshooting the System

The following topics address some system-level troubleshooting tasks and capabilities. For information on troubleshooting a specific feature, such as access control, see the chapter for the feature.

Pinging Addresses to Test Connectivity

Ping is a simple command that lets you determine if a particular address is alive and responsive. This means that basic connectivity is working. However, other policies running on a device could prevent specific types of traffic from successfully getting through a device. You can use **ping** by logging into the device CLI.



Note Because the system has multiple interfaces, you can control the interface used for pinging an address. You must ensure that you are using the right command, so that you are testing the connectivity that matters. For example, the system must be able to reach the Cisco license server through the virtual Management interface, so you must use the **ping system** command to test the connection. If you use **ping**, you are testing whether an address can be reached through the data interfaces, which might not give you the same result.

The normal ping uses ICMP packets to test the connection. If your network prohibits ICMP, you can use a TCP ping instead (for data interface pings only).

Following are the main options for pinging network addresses.

Pinging an address through the virtual Management interface

Use the **ping system** command.

ping system *host*

The host can be an IP address or fully-qualified domain name (FQDN), such as `www.example.com`. Unlike pings through the data interfaces, there is no default count for system pings. The ping continues until you stop it using Ctrl+c. For example:

```
> ping system www.cisco.com
PING origin-www.cisco.COM (72.163.4.161) 56(84) bytes of data.
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=1 ttl=242 time=10.6 ms
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=2 ttl=242 time=8.13 ms
```

```

64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=3 ttl=242 time=8.51 ms
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=4 ttl=242 time=8.40 ms
^C
--- origin-www.cisco.COM ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3003ms
rtt min/avg/max/mdev = 8.139/8.927/10.650/1.003 ms
>

```

Pinging an address through a data interface using the routing table

Use the **ping** command. Without specifying an interface, you are testing whether the system can generically find a route to the host. Because this is how the system normally routes traffic, this is typically what you want to test.

ping *host*

Specify the IP address of the host. If you only know the FQDN, use the **nslookup fqdn-name** command to determine the IP address. For example:

```

> ping 171.69.38.1
Sending 5, 100-byte ICMP Echos to 171.69.38.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/10 ms

```



Note You can specify the timeout, repeat count, packet size, and even the data pattern to send. Use the help indicator, **?**, in the CLI to see the available options.

Pinging an address through a specific data interface

Use the **ping interface if_name** command if you want to test connectivity through a specific data interface. You can also specify the diagnostic interface using this command, but not the virtual management interface.

ping interface *if_name* *host*

Specify the IP address of the host. If you only know the FQDN, use the **nslookup fqdn-name** command to determine the IP address. For example:

```

> ping interface inside 171.69.38.1
Sending 5, 100-byte ICMP Echos to 171.69.38.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/10 ms

```

Pinging an address through a data interface using TCP ping

Use the **ping tcp** command. A TCP ping sends SYN packets and considers the ping successful if the destination sends a SYN-ACK packet.

ping tcp [*interface if_name*] *host port*

You must specify the host and TCP port. If you only know the FQDN, use the **nslookup fqdn-name** command to determine the IP address.

You can optionally specify the interface, which is the source interface of the ping, not the interface through which to send the pings. This type of ping always uses the routing table.

A TCP ping sends SYN packets and considers the ping successful if the destination sends a SYN-ACK packet. For example:

```
> ping tcp 10.0.0.1 21
Type escape sequence to abort.
No source specified. Pinging from identity interface.
Sending 5 TCP SYN requests to 10.0.0.1 port 21
from 10.0.0.10, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```



Note You can also specify the timeout, repeat count, and the source address for the TCP ping. Use the help indicator, `?`, in the CLI to see the available options.

Tracing Routes to Hosts

If you are having problems sending traffic to an IP address, you can trace the route to the host to determine if there is a problem on the network path. A traceroute works by sending UDP packets on an invalid port, or ICMPv6 echoes, to a destination. The routers along the way to the destination respond with an ICMP Time Exceeded Message, and report that error to traceroute. Each node receives three packets, so you get three chances per node to get an informative result. You can use **traceroute** by logging into the device CLI.



Note There are separate commands for tracing a route through a data interface (**traceroute**) or through the virtual management interface (**traceroute system**). Ensure that you use the right command.

The following table describes the possible result per packet as displayed in the output.

Output Symbol	Description
*	No response was received for the probe within the timeout period.
<i>nn msec</i>	For each node, the round-trip time (in milliseconds) for the specified number of probes.
!N.	ICMP network unreachable.
!H	ICMP host unreachable.
!P	ICMP protocol unreachable.
!A	ICMP administratively prohibited.
?	Unknown ICMP error.

Tracing a route through the virtual management interface

Use the **traceroute system** command.

traceroute system *destination*

The host can be an IPv4/IPv6 address or fully-qualified domain name (FQDN), such as `www.example.com`. For example:

```
> traceroute system www.example.com
traceroute to www.example.com (172.163.4.161), 30 hops max, 60 byte packets
 1 192.168.0.254 (192.168.0.254) 0.213 ms 0.310 ms 0.328 ms
 2 10.88.127.1 (10.88.127.1) 0.677 ms 0.739 ms 0.899 ms
 3 lab-gw1.example.com (10.89.128.25) 0.638 ms 0.856 ms 0.864 ms
 4 04-bb-gw1.example.com (10.152.240.65) 1.169 ms 1.355 ms 1.409 ms
 5 wan-gw1.example.com (10.152.240.33) 0.712 ms 0.722 ms 0.790 ms
 6 wag-gw1.example.com (10.152.240.73) 13.868 ms 10.760 ms 11.187 ms
 7 rbb-gw2.example.com (172.30.4.85) 7.202 ms 7.301 ms 7.101 ms
 8 rbb-gw1.example.com (172.30.4.77) 8.162 ms 8.225 ms 8.373 ms
 9 sbb-gw1.example.com (172.16.16.210) 7.396 ms 7.548 ms 7.653 ms
10 corp-gw2.example.com (172.16.16.58) 7.413 ms 7.310 ms 7.431 ms
11 dmzbb-gw2.example.com (172.16.0.78) 7.835 ms 7.705 ms 7.702 ms
12 dmzdcc-gw2.example.com (172.16.0.190) 8.126 ms 8.193 ms 11.559 ms
13 dcz05n-gw1.example.com (172.16.2.106) 11.729 ms 11.728 ms 11.939 ms
14 www1.example.com (172.16.4.161) 11.645 ms 7.958 ms 7.936 ms
```

Tracing a route through a data interface

Use the `traceroute` command.

`traceroute destination`

Specify the IP address of the host. If you only know the FQDN, use the `nslookup fqdn-name` command to determine the IP address. For example:

```
> traceroute 209.165.200.225
Tracing the route to 209.165.200.225
 1 10.83.194.1 0 msec 10 msec 0 msec
 2 10.83.193.65 0 msec 0 msec 0 msec
 3 10.88.193.101 0 msec 10 msec 0 msec
 4 10.88.193.97 0 msec 0 msec 10 msec
 5 10.88.239.9 0 msec 10 msec 0 msec
 6 10.88.238.65 10 msec 10 msec 0 msec
 7 172.16.7.221 70 msec 70 msec 80 msec
 8 209.165.200.225 70 msec 70 msec 70 msec
```



Note You can specify the timeout, time to live, number of packets per node, and even the IP address or interface to use as the source of the traceroute. Use the help indicator, `?`, in the CLI to see the available options.

Troubleshooting NTP

The system relies on accurate and consistent time to function correctly and to ensure that events and other data points are handled accurately. You must configure at least one, but ideally three, Network Time Protocol (NTP) servers to ensure the system always has reliable time information.

The device summary connection diagram (click **Device** in the main menu) shows the status of the connection to the NTP server. If the status is yellow or orange, then there is an issue with the connection to the configured servers. If the connection problem persists (it is not just a momentary issue), try the following.

- First, ensure that you have at least three NTP servers configured on **Device > System Settings > NTP**. Although this is not a requirement, reliability is greatly enhanced if you have at least three NTP servers.
- Ensure that there is a network path between the management interface IP address (defined on **Device > System Settings > Management Interface**) and the NTP servers.
 - If the management interface gateway is the data interfaces, you can configure static routes to the NTP servers on **Device > Routing** if the default route is not adequate.
 - If you set an explicit management interface gateway, log into the device CLI and use the **ping system** command to test whether there is a network path to each NTP server.
- Log into the device CLI and check the status of the NTP servers with the following commands.
 - **show ntp**—This command shows basic information about the NTP servers and their availability. However, the connectivity status in the FDM uses additional information to indicate the status, so there can be inconsistency in what this command shows and what the connectivity status diagram shows.
 - **system support ntp**—This command includes the output of **show ntp** plus the output of the standard NTP command **ntpq**, which is documented with the NTP protocol. Use this command if you need to confirm NTP synchronization.

Look for the section “Results of ‘ntpq -pn.’” For example, you might see something like the following:

```
Results of 'ntpq -pn'
remote           : +216.229.0.50
refid            : 129.7.1.66
st              : 2
t               : u
when            : 704
poll            : 1024
reach           : 377
delay           : 90.455
offset          : 2.954
jitter         : 2.473
```

In this example, the + before the NTP server address indicates that it is a potential candidate. An asterisk here, *, indicates the current time source peer.

The NTP daemon (NTPD) uses a sliding window of eight samples from each one of the peers and picks out one sample, then the clock selection determines the true chimers and the false tickers. NTPD then determines the round-trip distance (the offset of a candidate must not be over one-half the round trip delay). If connection delays, packet loss, or server issues cause one or all the candidates to be rejected, you would see long delays in the synchronization. The adjustment also occurs over a very long period of time: the clock offset and oscillator errors must be resolved by the clock discipline algorithm and this can take hours.



Note If the refid is .LOCL., this indicates the peer is an undisciplined local clock, that is, it is using its local clock only to set the time. The FDM always marks the NTP connection yellow (not synchronized) if the selected peer is .LOCL. Normally, NTP does not select a .LOCL. candidate if a better one is available, which is why you should configure at least three servers.

Troubleshooting DNS for the Management Interface

You must configure at least one DNS server for use by the Management interface. The server is needed for cloud connections to services such as smart licensing, database updates (such as GeoDB, rules, and VDB), and any other activity that needs domain name resolution.

Configuring a DNS server is rather trivial. You simply enter the IP addresses of the DNS servers you use when you initially configure the device. You can later change them on the **Device > System Settings > DNS Server** page.

However, the system can fail to resolve fully-qualified domain names (FQDN) due to network connectivity issues or problems with the DNS server itself. If you find the system cannot use your DNS servers, consider the following actions to identify and resolve the problem.

Procedure

Step 1

Determine if you have a problem.

- a) Use SSH to log into the device CLI.
- b) Enter **ping system www.cisco.com**. If you get an “unknown host” message like the following, then the system could not resolve the domain name. If the ping is successful, then you are done: DNS is working. (Press Ctrl+C to stop the ping.)

```
> ping system www.cisco.com
ping: unknown host www.cisco.com
```

Note It is critical that you include the **system** keyword in the **ping** command. The **system** keyword sends the ping through the management IP address, which is the only interface that uses the management DNS server. Pinging **www.cisco.com** is also a good option, because you need a route to that server for smart licensing and updates.

Step 2

Verify the configuration for the management interface.

- a) Click **Device > System Settings > Management Interface**, and verify the following. If you make changes, the changes are applied immediately on clicking **Save**. If you change the Management address, you will need to reconnect and log back in.
 - The gateway IP address is correct for the Management network. If you using the data interfaces as the gateway, subsequent steps will verify that configuration.
 - If you are not using the data interfaces as a gateway, verify that the Management IP address/subnet mask and the gateway IP address are on the same subnet.
- b) Click **Device > System Settings > DNS Server** and verify that the right DNS servers are configured.

If you are deploying the device on your network edge, your service provider might have specific requirements about the DNS server you can use.
- c) If you are using the data interfaces as the gateway, verify that you have the required routes.

You need a default route for 0.0.0.0. You might need additional routes if the DNS server is not available through the gateway for the default route. There are two basic situations:

- If you are using DHCP to obtain an address for the outside interface, and you selected the **Obtain Default Route using DHCP** option, the default route is not visible in the FDM. From SSH, enter **show route** and verify that there is a route for 0.0.0.0. Because this is the default configuration for the outside interface, this is a likely situation that you might encounter. (Go to **Device > Interfaces** to view the configuration of the outside interface.)
- If you are using a static IP address on the outside interface, or you are not obtaining the default route from DHCP, then open **Device > Routing**. Verify that the correct gateway is being used for the default route.

If the DNS server cannot be reached through the default route, you must define a static route to it on the **Routing** page. Note that you should not add routes for directly connected networks, that is, networks that are connected directly to any of the system's data interfaces, as the system can route to those networks automatically.

Also verify that there are no static routes that are misdirecting traffic to the server out the wrong interface.

- d) If the deployment button indicates that there are undeployed changes, deploy them now and wait for deployment to complete.



- e) Retest **ping system www.cisco.com**. If you still have problems, continue with the next step.

Step 3

In the SSH session, enter **nslookup www.cisco.com**.

- If **nslookup** indicates that it got a response from the DNS server, but the server could not find the name, it means that DNS is configured correctly, but the DNS server you are using does not have an address for the FQDN. The response would look similar to the following:

```
> nslookup www.cisco.com
Server:      10.163.47.11
Address:    10.163.47.11#53

** server can't find www.cisco.com: NXDOMAIN
```

Resolution: In this case, you need to configure a different DNS server, or get the one you have updated so it can resolve the FQDNs you need resolved. Work with your network administrator or ISP to get the IP address of a DNS server that will work for your network.

- If you get a “connection timed out” message, then the system cannot reach your DNS servers, or all of the DNS servers are currently down and not responding (which is less likely). Continue with the next step.

```
> nslookup www.cisco.com
; ; connection timed out; no servers could be reached
```

Step 4

Use the **traceroute system DNS_server_ip_address** command to trace the route to the DNS server.

For example, if the DNS server is 10.100.10.1, enter:

```
> traceroute system 10.100.10.1
```

Following are the possible results:

- The traceroute completes and reaches the DNS server. In this case, there is in fact a route to the DNS server and the system can reach it. Thus, there is no routing problem. However, for some reason, DNS requests to this server are not getting a response.

Resolution: There is a possibility that a router or firewall along the path is dropping UDP/53 traffic, which is the port used for DNS. You might try a DNS server along a different network path. This is a difficult problem to resolve, as you will need to determine which node is blocking traffic, and work with the system administrator to get the access rules changed.

- The traceroute cannot reach even one node, which would look like the following:

```
> traceroute system 10.100.10.1
traceroute to 10.100.10.1 (10.100.10.1), 30 hops max, 60 byte packets
 1 * * *
 2 * * *
 3 * * *
 (and so forth)
```

Resolution: In this case, the routing problem is within your system. Try doing a **ping system** for the gateway IP address. Re-verify the configuration of the management interface as mentioned in earlier steps, and ensure that you have the required gateways and routes configured.

- The traceroute makes it through a few nodes before it can no longer resolve the route, which would look like the following:

```
> traceroute system 10.100.10.1
traceroute to 10.100.10.1 (10.100.10.1), 30 hops max, 60 byte packets
 1 192.168.0.254 (192.168.0.254) 0.475 ms 0.532 ms 0.542 ms
 2 10.88.127.1 (10.88.127.1) 0.803 ms 1.434 ms 1.443 ms
 3 site04-lab-gw1.example.com (10.89.128.25) 1.390 ms 1.399 ms 1.435 ms
 4 * * *
 5 * * *
 6 * * *
```

Resolution: In this case, routing breaks down at the last node. You might need to work with the system administrator to get correct routes installed in that node. However, if there is intentionally no route to the DNS server through the node, you need to change your gateway, or create your own static route, to point to a router that can route traffic to the DNS server.

Analyzing CPU and Memory Usage

To view system-level information about CPU and memory usage, select **Monitoring > System** and look for the CPU and Memory bar graphs. These graphs show information collected through the CLI using the **show cpu system** and **show memory system** commands.

If you log into the CLI, you can use additional versions of these commands to view other information. Typically, you would look at this information only if you are having persistent problems with usage, or at the direction of the Cisco Technical Assistance Center (TAC). Much of the detailed information is complex and requires TAC interpretation.

Following are some highlights of what you can examine. You can find more detailed information about these commands in [Cisco Firepower Threat Defense Command Reference](http://www.cisco.com/c/en/us/td/docs/security/firepower/command_ref/b_Command_Reference_for_Firepower_Threat_Defense.html) at http://www.cisco.com/c/en/us/td/docs/security/firepower/command_ref/b_Command_Reference_for_Firepower_Threat_Defense.html.

- **show cpu** displays data plane CPU utilization.
- **show cpu core** displays usage for each CPU core separately.
- **show cpu detailed** displays additional per-core and overall data plane CPU usage.
- **show memory** displays data plane memory usage.



Note Some of the keywords (not mentioned above) require that you first set up profiling or other features using the **cpu** or **memory** commands. Use these features at the direction of TAC only.

Viewing Logs

The system logs information for a wide variety of actions. You can use the **system support view-files** command to open a system log. Use this command while working with the Cisco Technical Assistance Center (TAC) so that they can help you interpret the output and to select the appropriate log to view.

The command presents a menu for selecting a log. Use the following commands to navigate the wizard:

- To change to a sub-directory, type in the name of the directory and press Enter.
- To select a file to view, enter **s** at the prompt. You are then prompted for a file name. You must type the complete name, and capitalization matters. The file list shows you the size of the log, which you might consider before opening very large logs.
- Press the space bar when you see --More-- to see the next page of log entries; press Enter to see just the next log entry. When you reach the end of the log, you are taken to the main menu. The --More-- line shows you the size of the log and how much of it you have viewed. **Use Ctrl+C to close the log and exit the command if you do not want to page through the entire log.**
- Type **b** to go up one level in the structure to the menu.

If you want to leave the log open so you can see new messages as they are added, use the **tail-logs** command instead of **system support view-files**.

The following example shows how view the `cisco/audit.log` file, which tracks attempts to log into the system. The file listing starts with directories at the top, then a list of files in the current directory.

```
> system support view-files

===View Logs===

=====
Directory: /ngfw/var/log
-----sub-dirs-----
cisco
mojo
removed_packages
setup
connector
```

```

sf
scripts
packages
removed_scripts
httpd
-----files-----
2016-10-14 18:12:04.514783 | 5371      | SMART_STATUS_sda.log
2016-10-14 18:12:04.524783 | 353      | SMART_STATUS_sdb.log
2016-10-11 21:32:23.848733 | 326517   | action_queue.log
2016-10-06 16:00:56.620019 | 1018     | brl.down.log

<list abbreviated>

([b] to go back or [s] to select a file to view, [Ctrl+C] to exit)
Type a sub-dir name to list its contents: cisco

=====
Directory: /ngfw/var/log/cisco
-----files-----
2017-02-13 22:44:42.394907 | 472      | audit.log
2017-02-13 23:40:30.858198 | 903615   | ev_stats.log.0
2017-02-09 18:14:26.870361 | 0        | ev_stats.log.0.lck
2017-02-13 05:24:00.682601 | 1024338  | ev_stats.log.1
2017-02-12 08:41:00.478103 | 1024338  | ev_stats.log.2
2017-02-11 11:58:00.260805 | 1024218  | ev_stats.log.3
2017-02-09 18:12:13.828607 | 95848    | firstboot.ngfw-onbox.log
2017-02-13 23:40:00.240359 | 6523160  | ngfw-onbox.log

([b] to go back or [s] to select a file to view, [Ctrl+C] to exit)
Type a sub-dir name to list its contents: s

Type the name of the file to view ([b] to go back, [Ctrl+C] to exit)
> audit.log
2017-02-09 18:59:26 - SubSystem:LOGIN, User:admin, IP:10.24.42.205, Message:Login successful,

2017-02-13 17:59:28 - SubSystem:LOGIN, User:admin, IP:10.24.111.72, Message:Login successful,

2017-02-13 22:44:36 - SubSystem:LOGIN, User:admin, IP:10.24.111.72, Message:Login failed,
2017-02-13 22:44:42 - SubSystem:LOGIN, User:admin, IP:10.24.111.72, Message:Login successful,

2017-02-13 22:44:42 - SubSystem:LOGIN, User:admin, IP:10.24.111.72, Message:Unlocked account.,

<remaining log truncated>

```

Creating a Troubleshooting File

Cisco Technical Assistance Center (TAC) personnel might ask you to submit system log information when you submit a problem report. This information assists them with diagnosing the problem. You do not need to submit a diagnostics file unless asked to do so.

The following procedure explains how to create and download the diagnostics file.

Procedure

Step 1 Click **Device**.

Step 2 Under **Troubleshooting**, click **Request File to be Created** or **Re-Request File to be Created** (if you have created one before).

The system starts generating the diagnostic file. You can go to other pages and return here to check status. When the file is ready, the date and time of the file creation is shown along with a download button.

Step 3 When the file is ready, click the download button.

The file is downloaded to your workstation using your browser's standard download method.

Uncommon Management Tasks

The following topics cover actions you would not perform often, if ever. All of these actions result in erasing your device configuration. Ensure that the device is not currently providing critical services to a production network before making these changes.

Switching Between Local and Remote Management

You can configure and manage your device using the local FDM, which is hosted directly on the device, or remotely, using the FMC multiple device manager. You might want to use the remote manager if you want to configure features not supported by the FDM, or if you need the power and analysis capabilities available in the FMC.

You also must use the FMC if you want to run the device in transparent firewall mode.

You can switch between local and remote management without reinstalling the software. Before switching from remote to local management, verify that the FDM meets all of your configuration requirements.



Caution Switching managers erases the device configuration and returns the system to the default configuration. However, management IP address and hostname are preserved.

Before you begin

If you registered the device, especially if you enabled any feature licenses, you must unregister the device through the FDM before switching to remote management. Unregistering the device frees the base license and all feature licenses. If you do not unregister the device, those licenses remain assigned to the device in Cisco Smart Software Manager. See [Unregistering the Device, on page 68](#).

Procedure

Step 1 Use an SSH client to open a connection to the **management IP address** and log into the device CLI with a username that has configuration CLI access. For example, the **admin** username.

It is important that you follow this process while connected to the management IP address. When using the FDM, you have the option to manage the device through the IP address on a data interface. However, you must use the Management physical port and management IP address to manage the device remotely.

If you cannot connect to the management IP address, address the following:

- Ensure that the Management physical port is wired to a functioning network.
- Ensure that the management IP address and gateway are configured for the management network. From the FDM, configure the address and gateway on **Device > System Settings > Management Interface**. (In the CLI, use the **configure network ipv4/ipv6 manual** command.)

Note Ensure that you are using an external gateway for the management IP address. You cannot use the data interfaces as a gateway when using a remote manager.

Step 2 To switch from local to remote management:

- a) Verify you are currently in local management mode.

```
> show managers
Managed locally.
```

- b) Configure the remote manager.

```
configure manager add {hostname | IPv4_address | IPv6_address | DONTRESOLVE} regkey [nat_id]
```

Where:

- {hostname | IPv4_address | IPv6_address | DONTRESOLVE} specifies the DNS host name or IP address (IPv4 or IPv6) of the FMC that manages this device. If the FMC is not directly addressable, use **DONTRESOLVE**. If you use **DONTRESOLVE**, *nat_id* is required.
- *regkey* is the unique alphanumeric registration key required to register a device to the FMC.
- *nat_id* is an optional alphanumeric string used during the registration process between the FMC and the device. It is required if the hostname is set to **DONTRESOLVE**.

For example, to use the manager at 192.168.0.123 with the registration key **secret**, enter the following:

```
> configure manager add 192.168.0.123 secret
If you enabled any feature licenses, you must disable them in
Firepower Device Manager before switching to remote management.
Otherwise, those licenses remain assigned to the device in Cisco
Smart Software Manager.
Do you want to continue [yes/no] yes
Manager successfully configured.
Please make note of reg_key as this will be required while adding
Device in FMC.

> show managers
Host                : 192.168.0.123
Registration Key     : ****
Registration         : pending
RPC Status           :
```

Note While registration is still pending, you can use **configure manager delete** to cancel the registration and then **configure manager local** to return to local management.

- c) Log into the FMC and add the device.

See the FMC online help for details.

Step 3 To switch from remote management to local management:

- a) Verify you are currently in remote management mode.

```
> show managers
Host           : 192.168.0.123
Registration Key : ****
Registration    : pending
RPC Status     :
```

- b) Delete the remote manager and go into no manager mode.

You cannot go directly from remote management to local management. Use the **configure manager delete** command to remove the manager.

```
> configure manager delete
Deleting task list
Manager successfully deleted.
```

```
>
> show managers
No managers configured.
```

- c) Configure the local manager.

configure manager local

For example:

```
> configure manager local
Deleting task list

> show managers
Managed locally.
```

You can now use a web browser to open the local manager at <https://management-IP-address>.

Changing the Firewall Mode

The FTD firewall can run in routed or transparent mode. A routed mode firewall is a routed hop and acts as a default gateway for hosts that connect to one of its screened subnets. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a "bump in the wire," or a "stealth firewall," and is not seen as a router hop to connected devices.

The local FDM supports routed mode only. If, however, you need to run the device in transparent mode, you can change the firewall mode and start managing the device with the FMC. Conversely, you can convert a transparent mode device to routed mode, and then you have the option to configure it with the local manager (you can also manage routed mode devices using FMC).

Regardless of local or remote management, you must use the device CLI to change the mode.

The following procedure explains how to change the mode when using the local manager, or when intending to use the local manager.



Caution Changing firewall mode erases the device configuration and returns the system to the default configuration. However, management IP address and hostname are preserved.

Before you begin

If you are converting to transparent mode, install the FMC before changing the firewall mode.

If you enabled any feature licenses, you must disable them in the FDM before deleting the local manager and switching to remote management. Otherwise, those licenses remain assigned to the device in Cisco Smart Software Manager. See [Enabling or Disabling Optional Licenses, on page 67](#).

Procedure

Step 1 Use an SSH client to open a connection to the **management IP address** and log into the device CLI with a username that has configuration CLI access. For example, the **admin** username.

It is important that you follow this process while connected to the management IP address. When using the FDM, you have the option to manage the device through the IP address on a data interface. However, you must use the Management physical port and management IP address to manage the device remotely.

If you cannot connect to the management IP address, address the following:

- Ensure that the Management physical port is wired to a functioning network.
- Ensure that the management IP address and gateway are configured for the management network. From the FDM, configure the address and gateway on **Device > System Settings > Management Interface**. (In the CLI, use the **configure network ipv4/ipv6 manual** command.)

Note Ensure that you are using an external gateway for the management IP address. You cannot use the data interfaces as a gateway when using a remote manager.

Step 2 To change the mode from routed to transparent and use remote management:

a) Disable local management and enter no manager mode.

You cannot change the firewall mode while there is an active manager. Use the **configure manager delete** command to remove the manager.

```
> configure manager delete
If you enabled any feature licenses, you must disable them in
Firepower Device Manager before deleting the local manager.
Otherwise, those licenses remain assigned to the device in
Cisco Smart Software Manager.
Do you want to continue[yes/no] yes
Deleting task list
Manager successfully deleted.

>
> show managers
No managers configured.
```

b) Change the firewall mode to transparent.

configure firewall transparent**Example:**

```
> configure firewall transparent
This will destroy the current interface configurations,
are you sure that you want to proceed? [y/N] y
The firewall mode was changed successfully.
```

- c) Configure the remote manager.

```
configure manager add {hostname | IPv4_address | IPv6_address | DONTRESOLVE} regkey [nat_id]
```

Where:

- {hostname | IPv4_address | IPv6_address | DONTRESOLVE} specifies the DNS host name or IP address (IPv4 or IPv6) of the FMC that manages this device. If the FMC is not directly addressable, use **DONTRESOLVE**. If you use **DONTRESOLVE**, *nat_id* is required.
- *regkey* is the unique alphanumeric registration key required to register a device to the FMC.
- *nat_id* is an optional alphanumeric string used during the registration process between the FMC and the device. It is required if the hostname is set to **DONTRESOLVE**.

For example, to use the manager at 192.168.0.123 with the registration key **secret**, enter the following:

```
> configure manager add 192.168.0.123 secret
Manager successfully configured.
Please make note of reg_key as this will be required while adding
Device in FMC.

> show managers
Host                : 192.168.0.123
Registration Key    : ****
Registration        : pending
RPC Status          :
```

- d) Log into the FMC and add the device.

See the FMC online help for details.

Step 3

To change the mode from transparent to routed and convert to local management:

- Deregister the device from the FMC.
- Access the FTD device CLI, preferably from the console port.

Because changing the mode erases your configuration, the management IP address will revert to the default, so you might lose an SSH connection to the management IP address after changing modes.

- c) Change the firewall mode to routed.

configure firewall routed**Example:**

```
> configure firewall routed
This will destroy the current interface configurations,
are you sure that you want to proceed? [y/N] y
The firewall mode was changed successfully.
```

- d) Enable the local manager.

configure manager local

For example:

```
> configure manager local
Deleting task list

> show managers
Managed locally.
```

You can now use a web browser to open the local manager at **https://management-IP-address**.

Resetting the Configuration

You can reset the system configuration to the factory default if you want to start over. Although you cannot directly reset the configuration, deleting and adding the manager clears the configuration.

If your intention is to wipe away the configuration and then recover a backup, ensure that you have already download the backup copy you want to restore. You will need to upload it after resetting the system so that you can restore it.

Before you begin

If you enabled any feature licenses, you must disable them in the FDM before deleting the local manager. Otherwise, those licenses remain assigned to the device in Cisco Smart Software Manager. See [Enabling or Disabling Optional Licenses, on page 67](#).

Procedure

Step 1 Use an SSH client to open a connection to the management IP address and log into the device CLI with a username that has configuration CLI access. For example, the **admin** username.

Step 2 Use the **configure manager delete** command to remove the manager.

```
> configure manager delete
If you enabled any feature licenses, you must disable them in
Device Manager before deleting the local manager.
Otherwise, those licenses remain assigned to the device in Cisco
Smart Software Manager.
Do you want to continue[yes/no] yes
Deleting task list
Manager successfully deleted.

>
> show managers
No managers configured.
```

Step 3 Configure the local manager.
configure manager local

For example:

```
> configure manager local  
Deleting task list
```

```
> show managers  
Managed locally.
```

You can now use a web browser to open the local manager at <https://management-IP-address>. By clearing the configuration, you will be prompted to complete the device setup wizard.
