Cisco ASA Series Firewall ASDM Configuration Guide

Software Version 7.2

For the ASA 5505, ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, ASA 5555-X, ASA 5585-X, ASA Services Module, and the Adaptive Security Virtual Appliance

Released: April 24, 2014
Updated: September 16, 2014

Cisco Systems, Inc.

www.cisco.com

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

CONTENTS

CHAPTER 7   **Twice NAT (ASA 8.3 and Later)**   **7-1**

PART 3   **Application Inspection**

CHAPTER 8   **Getting Started with Application Layer Protocol Inspection**   **8-1**

PART 4   **Unified Communications**

CHAPTER 13   **Information About the ASA in Cisco Unified Communications**   **13-1**

CHAPTER 14   **Cisco Unified Communication Wizard**   **14-1**

PART 6   **Advanced Network Protection**

CHAPTER 23   **ASA and Cisco Cloud Web Security**   **23-1**

PART 7   **ASA Modules**

CHAPTER 26   **ASA FirePOWER (SFR) Module**   **26-1**

About This Guide

- Document Objectives, page xxi
- Related Documentation, page xxi
- Conventions, page xxii
- Obtaining Documentation and Submitting a Service Request, page xxii

Document Objectives

The purpose of this guide is to help you configure the firewall features for Cisco ASA series using the Adaptive Security Device Manager (ASDM). This guide does not cover every feature, but describes only the most common configuration scenarios.

Throughout this guide, the term "ASA" applies generically to supported models, unless specified otherwise.

Note ASDM supports many ASA versions. The ASDM documentation and online help includes all of the latest features supported by the ASA. If you are running an older version of ASA software, the documentation might include features that are not supported in your version. Similarly, if a feature was added into a maintenance release for an older major or minor version, then the ASDM documentation includes the new feature even though that feature might not be available in all later ASA releases. Please refer to the feature history table for each chapter to determine when features were added. For the minimum supported version of ASDM for each ASA version, see *Cisco ASA Series Compatibility*.

Related Documentation

For more information, see *Navigating the Cisco ASA Series Documentation* at http://www.cisco.com/go/asadocs.

Conventions

This document uses the following conventions:

Convention	Indication
bold font	Commands and keywords and user-entered text appear in **bold** font.
italic font	Document titles, new or emphasized terms, and arguments for which you supply values are in *italic* font.
[ ]	Elements in square brackets are optional.
{x | y | z }	Required alternative keywords are grouped in braces and separated by vertical bars.
[x | y | z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.
`courier` font	Terminal sessions and information the system displays appear in `courier` font.
`courier bold` font	Commands and keywords and user-entered text appear in **`bold courier`** font.
`courier italic` font	Arguments for which you supply values are in *`courier italic`* font.
< >	Nonprinting characters such as passwords are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

Note Means *reader take note*.

Tip Means *the following information will help you solve a problem*.

Caution Means *reader be careful*. In this situation, you might perform an action that could result in equipment damage or loss of data.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see *What's New in Cisco Product Documentation* at: http://www.cisco.com/c/en/us/td/docs/general/whatsnew/whatsnew.html.

Subscribe to *What's New in Cisco Product Documentation*, which lists all new and revised Cisco technical documentation as an RSS feed and delivers content directly to your desktop using a reader application. The RSS feeds are a free service.

P ART 1

Service Policies and Access Control

Service Policy

Released: April 24, 2014
Updated: September 16, 2014

Service policies provide a consistent and flexible way to configure ASA features. For example, you can use a service policy to create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications. A service policy consists of multiple service policy rules applied to an interface or applied globally.

This chapter includes the following sections:

- Information About Service Policies, page 1-1
- Licensing Requirements for Service Policies, page 1-5
- Guidelines and Limitations, page 1-6
- Default Settings, page 1-7
- Task Flows for Configuring Service Policies, page 1-8
- Adding a Service Policy Rule for Through Traffic, page 1-8
- Adding a Service Policy Rule for Management Traffic, page 1-13
- Managing the Order of Service Policy Rules, page 1-15
- Feature History for Service Policies, page 1-17

Information About Service Policies

This section describes how service policies work and includes the following topics:

- Supported Features, page 1-2
- Feature Directionality, page 1-2
- Feature Matching Within a Service Policy, page 1-3
- Order in Which Multiple Feature Actions are Applied, page 1-4
- Incompatibility of Certain Feature Actions, page 1-5
- Feature Matching for Multiple Service Policies, page 1-5

Supported Features

Table 1-1 lists the features supported by service policy rules.

Table 1-1   ***Service Policy Rule Features***

Feature	For Through Traffic?	For Management Traffic?	See:
Application inspection (multiple types)	**All** except RADIUS accounting	**RADIUS accounting** only	• Chapter 8, "Getting Started with Application Layer Protocol Inspection." • Chapter 9, "Inspection of Basic Internet Protocols." • Chapter 10, "Inspection for Voice and Video Protocols." • Chapter 11, "Inspection of Database and Directory Protocols." • Chapter 12, "Inspection for Management Application Protocols." • Chapter 23, "ASA and Cisco Cloud Web Security."
ASA IPS	**Yes**	No	Chapter 28, "ASA IPS Module."
ASA CX	**Yes**	No	Chapter 27, "ASA CX Module."
ASA FirePOWER (ASA SFR)	**Yes**	No	Chapter 26, "ASA FirePOWER (SFR) Module."
NetFlow Secure Event Logging filtering	**Yes**	**Yes**	See the general operations configuration guide.
QoS input and output policing	**Yes**	No	Chapter 21, "Quality of Service."
QoS standard priority queue	**Yes**	No	Chapter 21, "Quality of Service."
QoS traffic shaping, hierarchical priority queue	**Yes**	**Yes**	Chapter 21, "Quality of Service."
TCP and UDP connection limits and timeouts, and TCP sequence number randomization	**Yes**	**Yes**	Chapter 20, "Connection Settings."
TCP normalization	**Yes**	No	Chapter 20, "Connection Settings."
TCP state bypass	**Yes**	No	Chapter 20, "Connection Settings."
User statistics for Identity Firewall	**Yes**	**Yes**	See the **user-statistics** command in the command reference.

Feature Directionality

Actions are applied to traffic bidirectionally or unidirectionally depending on the feature. For features that are applied bidirectionally, all traffic that enters or exits the interface to which you apply the policy map is affected if the traffic matches the class map for both directions.

> **Note** When you use a global policy, all features are unidirectional; features that are normally bidirectional when applied to a single interface only apply to the ingress of each interface when applied globally. Because the policy is applied to all interfaces, the policy will be applied in both directions so bidirectionality in this case is redundant.

For features that are applied unidirectionally, for example QoS priority queue, only traffic that enters (or exits, depending on the feature) the interface to which you apply the policy map is affected. See Table 1-2 for the directionality of each feature.

Table 1-2   Feature Directionality

Feature	Single Interface Direction	Global Direction
Application inspection (multiple types)	Bidirectional	Ingress
ASA CSC	Bidirectional	Ingress
ASA CX	Bidirectional	Ingress
ASA CX authentication proxy	Ingress	Ingress
ASA FirePOWER (ASA SFR)	Bidirectional	Ingress
ASA IPS	Bidirectional	Ingress
NetFlow Secure Event Logging filtering	N/A	Ingress
QoS input policing	Ingress	Ingress
QoS output policing	Egress	Egress
QoS standard priority queue	Egress	Egress
QoS traffic shaping, hierarchical priority queue	Egress	Egress
TCP and UDP connection limits and timeouts, and TCP sequence number randomization	Bidirectional	Ingress
TCP normalization	Bidirectional	Ingress
TCP state bypass	Bidirectional	Ingress
User statistics for Identity Firewall	Bidirectional	Ingress

Feature Matching Within a Service Policy

See the following information for how a packet matches rules in a policy for a given interface:

1. A packet can match only one rule for an interface for each feature type.

2. When the packet matches a rule for a feature type, the ASA does not attempt to match it to any subsequent rules for that feature type.

3. If the packet matches a subsequent rule for a different feature type, however, then the ASA also applies the actions for the subsequent rule, if supported. See Incompatibility of Certain Feature Actions, page 1-5 for more information about unsupported combinations.

> **Note** Application inspection includes multiple inspection types, and most are mutually exclusive. For inspections that can be combined, each inspection is considered to be a separate feature.

For example, if a packet matches a rule for connection limits, and also matches a rule for an application inspection, then both actions are applied.

If a packet matches a rulefor HTTP inspection, but also matches another rule that includes HTTP inspection, then the second rule actions are not applied.

If a packet matches a rulefor HTTP inspection, but also matches another rule that includes FTP inspection, then the second rule actions are not applied because HTTP and FTP inspections cannpt be combined.

If a packet matches a rule for HTTP inspection, but also matches another rule that includes IPv6 inspection, then both actions are applied because the IPv6 inspection can be combined with any other type of inspection.

Order in Which Multiple Feature Actions are Applied

The order in which different types of actions in a service policy are performed is independent of the order in which the actions appear in the table.

Note   NetFlow Secure Event Logging filtering and User statistics for Identity Firewall are order-independent.

Actions are performed in the following order:

1. QoS input policing

2. TCP normalization, TCP and UDP connection limits and timeouts, TCP sequence number randomization, and TCP state bypass.

Note   When a the ASA performs a proxy service (such as AAA or CSC) or it modifies the TCP payload (such as FTP inspection), the TCP normalizer acts in dual mode, where it is applied before and after the proxy or payload modifying service.

3. ASA CSC

4. Application inspections that can be combined with other inspections:

    a. IPv6

    b. IP options

    c. WAAS

5. Application inspections that cannot be combined with other inspections. See Incompatibility of Certain Feature Actions, page 1-5 for more information.

6. ASA IPS

7. ASA CX

8. ASA FirePOWER (ASA SFR)

9. QoS output policing

10. QoS standard priority queue

11. QoS traffic shaping, hierarchical priority queue

Incompatibility of Certain Feature Actions

Some features are not compatible with each other for the same traffic. The following list may not include all incompatibilities; for information about compatibility of each feature, see the chapter or section for your feature:

- You cannot configure QoS priority queueing and QoS policing for the same set of traffic.

- Most inspections should not be combined with another inspection, so the ASA only applies one inspection if you configure multiple inspections for the same traffic. HTTP inspection can be combined with the Cloud Web Security inspection. Other exceptions are listed in the Order in Which Multiple Feature Actions are Applied, page 1-4.

- You cannot configure traffic to be sent to multiple modules, such as the ASA CX and ASA IPS.

- HTTP inspection is not compatible with the ASA CX or the ASA FirePOWER.

- The ASA CX and ASA FirePOWER modules are not compatible with Cloud Web Security.

Note   The Default Inspection Traffic traffic class, which is used in the default global policy, is a special CLI shortcut to match the default ports for all inspections. When used in a policy map, this class map ensures that the correct inspection is applied to each packet, based on the destination port of the traffic. For example, when UDP traffic for port 69 reaches the ASA, then the ASA applies the TFTP inspection; when TCP traffic for port 21 arrives, then the ASA applies the FTP inspection. So in this case only, you can configure multiple inspections for the same class map. Normally, the ASA does not use the port number to determine which inspection to apply, thus giving you the flexibility to apply inspections to non-standard ports, for example.

This traffic class does not include the default ports for Cloud Web Security inspection (80 and 443).

Feature Matching for Multiple Service Policies

For TCP and UDP traffic (and ICMP when you enable stateful ICMP inspection), service policies operate on traffic flows, and not just individual packets. If traffic is part of an existing connection that matches a feature in a policy on one interface, that traffic flow cannot also match the same feature in a policy on another interface; only the first policy is used.

For example, if HTTP traffic matches a policy on the inside interface to inspect HTTP traffic, and you have a separate policy on the outside interface for HTTP inspection, then that traffic is not also inspected on the egress of the outside interface. Similarly, the return traffic for that connection will not be inspected by the ingress policy of the outside interface, nor by the egress policy of the inside interface.

For traffic that is not treated as a flow, for example ICMP when you do not enable stateful ICMP inspection, returning traffic can match a different policy map on the returning interface. For example, if you configure IPS on the inside and outside interfaces, but the inside policy uses virtual sensor 1 while the outside policy uses virtual sensor 2, then a non-stateful Ping will match virtual sensor 1 outbound, but will match virtual sensor 2 inbound.

Licensing Requirements for Service Policies

Model	License Requirement
ASAv	Standard or Premium License.
All other models	Base License.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

IPv6 Guidelines

Supports IPv6 for the following features:

- Application inspection for DNS, FTP, HTTP, ICMP, ScanSafe, SIP, SMTP, IPsec-pass-thru, and IPv6.
- ASA IPS
- ASA CX
- ASA FirePOWER
- NetFlow Secure Event Logging filtering
- TCP and UDP connection limits and timeouts, TCP sequence number randomization
- TCP normalization
- TCP state bypass
- User statistics for Identity Firewall

Traffic Class Guidelines

The maximum number of traffic classes of all types is 255 in single mode or per context in multiple mode. Class maps include the following types:

- Layer 3/4 class maps (for through traffic and management traffic).
- Inspection class maps
- Regular expression class maps
- **match** commands used directly underneath an inspection policy map

This limit also includes default traffic classes of all types, limiting user-configured traffic classes to approximately 235. See Default Traffic Classes, page 1-8.

Service Policy Guidelines

- Interface service policies take precedence over the global service policy for a given feature. For example, if you have a global policy with FTP inspection, and an interface policy with TCP normalization, then both FTP inspection and TCP normalization are applied to the interface. However, if you have a global policy with FTP inspection, and an interface policy with FTP inspection, then only the interface policy FTP inspection is applied to that interface.

- You can only apply one global policy. For example, you cannot create a global policy that includes feature set 1, and a separate global policy that includes feature set 2. All features must be included in a single policy.

- When you make service policy changes to the configuration, all *new* connections use the new service policy. Existing connections continue to use the policy that was configured at the time of the connection establishment. **show** command output will not include data about the old connections.

  For example, if you remove a QoS service policy from an interface, then re-add a modified version, then the **show service-policy** command only displays QoS counters associated with new connections that match the new service policy; existing connections on the old policy no longer show in the command output.

  To ensure that all connections use the new policy, you need to disconnect the current connections so they can reconnect using the new policy. See the **clear conn** or **clear local-host** commands.

Default Settings

The following topics describe the default settings for Modular Policy Framework:

- Default Configuration, page 1-7
- Default Traffic Classes, page 1-8

Default Configuration

By default, the configuration includes a policy that matches all default application inspection traffic and applies certain inspections to the traffic on all interfaces (a global policy). Not all inspections are enabled by default. You can only apply one global policy, so if you want to alter the global policy, you need to either edit the default policy or disable it and apply a new one. (An interface policy overrides the global policy for a particular feature.)

The default policy includes the following application inspections:

- DNS
- FTP
- H323 (H225)
- H323 (RAS)
- RSH
- RTSP
- ESMTP
- SQLnet
- Skinny (SCCP)
- SunRPC

- XDMCP
- SIP
- NetBios
- TFTP
- IP Options

Default Traffic Classes

The configuration includes a default traffic class that the ASA uses in the default global policy called Default Inspection Traffic; it matches the default inspection traffic. This class, which is used in the default global policy, is a special shortcut to match the default ports for all inspections. When used in a policy, this class ensures that the correct inspection is applied to each packet, based on the destination port of the traffic. For example, when UDP traffic for port 69 reaches the ASA, then the ASA applies the TFTP inspection; when TCP traffic for port 21 arrives, then the ASA applies the FTP inspection. So in this case only, you can configure multiple inspections for the same class map. Normally, the ASA does not use the port number to determine which inspection to apply, thus giving you the flexibility to apply inspections to non-standard ports, for example.

Another class map that exists in the default configuration is called class-default, and it matches all traffic. You can use the class-default class if desired, rather than using the Any traffic class. In fact, some features are only available for class-default, such as QoS traffic shaping.

Task Flows for Configuring Service Policies

This section includes the following topics:

- Task Flow for Configuring a Service Policy Rule, page 1-8

Task Flow for Configuring a Service Policy Rule

Configuring a service policy consists of adding one or more service policy rules per interface or for the global policy. For each rule, you identify the following elements:

Step 1   Identify the interface to which you want to apply the rule, or identify the global policy.

Step 2   Identify the traffic to which you want to apply actions. You can identify Layer 3 and 4 through traffic.

Step 3   Apply actions to the traffic class. You can apply multiple actions for each traffic class.

Adding a Service Policy Rule for Through Traffic

See Supported Features, page 1-2 for more information. To add a service policy rule for through traffic, perform the following steps:

Step 1 Choose **Configuration** > **Firewall** > **Service Policy Rules** pane, and click **Add**.

The Add Service Policy Rule Wizard - Service Policy dialog box appears.

✎
Note When you click the Add button, and not the small arrow on the right of the Add button, you add a through traffic rule by default. If you click the arrow on the Add button, you can choose between a through traffic rule and a management traffic rule.

Step 2 In the Create a Service Policy and Apply To area, click one of the following options:

- **Interface**. This option applies the service policy to a single interface. Interface service policies take precedence over the global service policy for a given feature. For example, if you have a global policy with FTP inspection, and an interface policy with TCP connection limits, then both FTP inspection and TCP connection limits are applied to the interface. However, if you have a global policy with FTP inspection, and an interface policy with FTP inspection, then only the interface policy FTP inspection is applied to that interface.

  a. Choose an interface from the drop-down list.

     If you choose an interface that already has a policy, then the wizard lets you add a new service policy rule to the interface.

  b. If it is a new service policy, enter a name in the Policy Name field.

  c. (Optional) Enter a description in the Description field.

d. (Optional) Check the **Drop and log unsupported IPv6 to IPv6 traffic** check box to generate a syslog (767001) for IPv6 traffic that is dropped by application inspections that do not support IPv6 traffic. By default, syslogs are not generated. For a list of inspections that support IPv6, see IPv6 Guidelines, page 1-6.

- **Global - applies to all interfaces**. This option applies the service policy globally to all interfaces. By default, a global policy exists that includes a service policy rule for default application inspection. See Default Settings, page 1-7 for more information. You can add a rule to the global policy using the wizard.

  a. If it is a new service policy, enter a name in the Policy Name field.

  b. (Optional) Enter a description in the Description field.

  c. (Optional) Check the **Drop and log unsupported IPv6 to IPv6 traffic** check box to generate a syslog (767001) for IPv6 traffic that is dropped by application inspections that do not support IPv6 traffic. By default, syslogs are not generated. For a list of inspections that support IPv6, see IPv6 Guidelines, page 1-6.

Step 3 Click **Next**.

The Add Service Policy Rule Wizard - Traffic Classification Criteria dialog box appears.

Step 4 Click one of the following options to specify the traffic to which to apply the policy actions:

- **Create a new traffic class**. Enter a traffic class name in the Create a new traffic class field, and enter an optional description.

  Identify the traffic using one of several criteria:

  - **Default Inspection Traffic**—The class matches the default TCP and UDP ports used by all applications that the ASA can inspect.

    This option, which is used in the default global policy, is a special shortcut that when used in a rule, ensures that the correct inspection is applied to each packet, based on the destination port of the traffic. For example, when UDP traffic for port 69 reaches the ASA, then the ASA applies the TFTP inspection; when TCP traffic for port 21 arrives, then the ASA applies the FTP inspection. So in this case only, you can configure multiple inspections for the same rule (See Incompatibility of Certain Feature Actions, page 1-5 for more information about combining actions). Normally, the ASA does not use the port number to determine the inspection applied, thus giving you the flexibility to apply inspections to non-standard ports, for example.

    See Default Settings and NAT Limitations, page 8-4 for a list of default ports. The ASA includes a default global policy that matches the default inspection traffic, and applies common inspections to the traffic on all interfaces. Not all applications whose ports are included in the Default Inspection Traffic class are enabled by default in the policy map.

    You can specify a Source and Destination IP Address (uses ACL) class along with the Default Inspection Traffic class to narrow the matched traffic. Because the Default Inspection Traffic class specifies the ports and protocols to match, any ports and protocols in the ACL are ignored.

  - **Source and Destination IP Address (uses ACL)**—The class matches traffic specified by an extended ACL. If the ASA is operating in transparent firewall mode, you can use an EtherType ACL.

    Note When you create a new traffic class of this type, you can only specify one access control entry (ACE) initially. After you finish adding the rule, you can add additional ACEs by adding a new rule to the same interface or global policy, and then specifying **Add rule to existing traffic class** on the Traffic Classification dialog box (see below).

- **Tunnel Group**—The class matches traffic for a tunnel group to which you want to apply QoS. You can also specify one other traffic match option to refine the traffic match, excluding Any Traffic, Source and Destination IP Address (uses ACL), or Default Inspection Traffic.

- **TCP or UDP Destination Port**—The class matches a single port or a contiguous range of ports.

Tip   For applications that use multiple, non-contiguous ports, use the Source and Destination IP Address (uses ACL) to match each port.

- **RTP Range**—The class map matches RTP traffic.

- **IP DiffServ CodePoints (DSCP)**—The class matches up to eight DSCP values in the IP header.

- **IP Precedence**—The class map matches up to four precedence values, represented by the TOS byte in the IP header.

- **Any Traffic**—Matches all traffic.

- **Add rule to existing traffic class**. If you already have a service policy rule on the same interface, or you are adding to the global service policy, this option lets you add an ACE to an existing ACL. You can add an ACE to any ACL that you previously created when you chose the Source and Destination IP Address (uses ACL) option for a service policy rule on this interface. For this traffic class, you can have only one set of rule actions even if you add multiple ACEs. You can add multiple ACEs to the same traffic class by repeating this entire procedure. See Managing the Order of Service Policy Rules, page 1-15 for information about changing the order of ACEs.

- **Use an existing traffic class.** If you created a traffic class used by a rule on a different interface, you can reuse the traffic class definition for this rule. Note that if you alter the traffic class for one rule, the change is inherited by all rules that use that traffic class. If your configuration includes any **class-map** commands that you entered at the CLI, those traffic class names are also available (although to view the definition of the traffic class, you need to create the rule).

- **Use class default as the traffic class**. This option uses the class-default class, which matches all traffic. The class-default class is created automatically by the ASA and placed at the end of the policy. If you do not apply any actions to it, it is still created by the ASA, but for internal purposes only. You can apply actions to this class, if desired, which might be more convenient than creating a new traffic class that matches all traffic. You can only create one rule for this service policy using the class-default class, because each traffic class can only be associated with a single rule per service policy.

Step 5   Click **Next**.

Step 6   The next dialog box depends on the traffic match criteria you chose.

Note   The Any Traffic option does not have a special dialog box for additional configuration.

- Default Inspections—This dialog box is informational only, and shows the applications and the ports that are included in the traffic class.

- Source and Destination Address—This dialog box lets you set the source and destination addresses:

    a. Click **Match** or **Do Not Match**.

    The Match option creates a rule where traffic matching the addresses have actions applied. The Do Not Match option exempts the traffic from having the specified actions applied. For example, you want to match all traffic in 10.1.1.0/24 and apply connection limits to it, except

for 10.1.1.25. In this case, create two rules, one for 10.1.1.0/24 using the Match option and one for 10.1.1.25 using the Do Not Match option. Be sure to arrange the rules so that the Do Not Match rule is above the Match rule, or else 10.1.1.25 will match the Match rule first.

b. In the Source field, enter the source IP address, or click the **...** button to choose an IP address that you already defined in ASDM.

Specify the address and subnet mask using prefix/length notation, such as 10.1.1.0/24. If you enter an IP address without a mask, it is considered to be a host address, even if it ends with a 0.

Enter **any** to specify any source address.

Separate multiple addresses by a comma.

c. In the Destination field, enter the destination IP address, or click the **...** button to choose an IP address that you already defined in ASDM.

Specify the address and subnet mask using prefix/length notation, such as 10.1.1.0/24. If you enter an IP address without a mask, it is considered to be a host address, even if it ends with a 0.

Enter **any** to specify any destination address.

Separate multiple addresses by a comma.

d. In the Service field, enter an IP service name or number for the destination service, or click the **...** button to choose a service.

If you want to specify a TCP or UDP port number, or an ICMP service number, enter *protocol/port*. For example, enter TCP/8080.

By default, the service is IP.

Separate multiple services by a comma.

e. (Optional) Enter a description in the Description field.

f. (Optional) To specify a source service for TCP or UDP, click the **More Options** area open, and enter a TCP or UDP service in the Source Service field.

The destination service and source service must be the same. Copy and paste the destination Service field to the Source Service field.

g. (Optional) To make the rule inactive, click the **More Options** area open, and uncheck **Enable Rule**.

This setting might be useful if you do not want to remove the rule, but want to turn it off.

h. (Optional) To set a time range for the rule, click the **More Options** area open, and from the Time Range drop-down list, choose a time range.

To add a new time range, click the **...** button. See the general operations configuration guide for more information.

This setting might be useful if you only want the rule to be active at predefined times.

- Tunnel Group—Choose a tunnel group from the Tunnel Group drop-down list, or click **New** to add a new tunnel group. See the VPN configuration guide for more information.

To police each flow, check **Match flow destination IP address**. All traffic going to a unique IP destination address is considered a flow.

- Destination Port—Click **TCP** or **UDP**.

In the Service field, enter a port number or name, or click **...** to choose one already defined in ASDM.

- RTP Range—Enter an RTP port range, between 2000 and 65534. The maximum number of port sin the range is 16383.

- IP DiffServ CodePoints (DSCP)—In the DSCP Value to Add area, choose a value from the **Select Named DSCP Values** or enter a value in the **Enter DSCP Value (0-63) field**, and click **Add**.

  Add additional values as desired, or remove them using the **Remove** button.

- IP Precedence—From the Available IP Precedence area, choose a value and click **Add**.

  Add additional values as desired, or remove them using the **Remove** button.

Step 7   Click **Next**.

The Add Service Policy Rule - Rule Actions dialog box appears.

Step 8   Configure one or more rule actions. See Supported Features, page 1-2 for a list of features.

Step 9   Click **Finish**.

Adding a Service Policy Rule for Management Traffic

You can create a service policy for traffic directed to the ASA for management purposes. See Supported Features, page 1-2 for more information. This section includes the following topics:

Configuring a Service Policy Rule for Management Traffic

To add a service policy rule for management traffic, perform the following steps:

Step 1   From the Configuration > Firewall > Service Policy Rules pane, click the down arrow next to Add.

Step 2   Choose **Add Management Service Policy Rule**.

The Add Management Service Policy Rule Wizard - Service Policy dialog box appears.

Step 3   In the Create a Service Policy and Apply To area, click one of the following options:

- **Interface**. This option applies the service policy to a single interface. Interface service policies take precedence over the global service policy for a given feature. For example, if you have a global policy with RADIUS accounting inspection, and an interface policy with connection limits, then both RADIUS accounting and connection limits are applied to the interface. However, if you have a global policy with RADIUS accounting, and an interface policy with RADIUS accounting, then only the interface policy RADIUS accounting is applied to that interface.

  a. Choose an interface from the drop-down list.

    If you choose an interface that already has a policy, then the wizard lets you add a new service policy rule to the interface.

  b. If it is a new service policy, enter a name in the Policy Name field.

  c. (Optional) Enter a description in the Description field.

- **Global - applies to all interfaces**. This option applies the service policy globally to all interfaces. By default, a global policy exists that includes a service policy rule for default application inspection. See Default Settings, page 1-7 for more information. You can add a rule to the global policy using the wizard.

Step 4   Click **Next**.

The Add Management Service Policy Rule Wizard - Traffic Classification Criteria dialog box appears.

Step 5   Click one of the following options to specify the traffic to which to apply the policy actions:

- **Create a new traffic class**. Enter a traffic class name in the Create a new traffic class field, and enter an optional description.

   Identify the traffic using one of several criteria:

   - **Source and Destination IP Address (uses ACL)**—The class matches traffic specified by an extended ACL. If the ASA is operating in transparent firewall mode, you can use an EtherType ACL.

      ✎

      Note   When you create a new traffic class of this type, you can only specify one access control entry (ACE) initially. After you finish adding the rule, you can add additional ACEs by adding a new rule to the same interface or global policy, and then specifying **Add rule to existing traffic class** on the Traffic Classification dialog box (see below).

   - **TCP or UDP Destination Port**—The class matches a single port or a contiguous range of ports.

   🔍

   Tip   For applications that use multiple, non-contiguous ports, use the Source and Destination IP Address (uses ACL) to match each port.

- **Add rule to existing traffic class**. If you already have a service policy rule on the same interface, or you are adding to the global service policy, this option lets you add an ACE to an existing ACL. You can add an ACE to any ACL that you previously created when you chose the Source and Destination IP Address (uses ACL) option for a service policy rule on this interface. For this traffic class, you can have only one set of rule actions even if you add multiple ACEs. You can add multiple ACEs to the same traffic class by repeating this entire procedure. See Managing the Order of Service Policy Rules, page 1-15 for information about changing the order of ACEs.

- **Use an existing traffic class.** If you created a traffic class used by a rule on a different interface, you can reuse the traffic class definition for this rule. Note that if you alter the traffic class for one rule, the change is inherited by all rules that use that traffic class. If your configuration includes any **class-map** commands that you entered at the CLI, those traffic class names are also available (although to view the definition of the traffic class, you need to create the rule).

Step 6   Click **Next**.

Step 7   The next dialog box depends on the traffic match criteria you chose.

- Source and Destination Address—This dialog box lets you set the source and destination addresses:

   a. Click **Match** or **Do Not Match**.

      The Match option creates a rule where traffic matching the addresses have actions applied. The Do Not Match option exempts the traffic from having the specified actions applied. For example, you want to match all traffic in 10.1.1.0/24 and apply connection limits to it, except for 10.1.1.25. In this case, create two rules, one for 10.1.1.0/24 using the Match option and one for 10.1.1.25 using the Do Not Match option. Be sure to arrange the rules so that the Do Not Match rule is above the Match rule, or else 10.1.1.25 will match the Match rule first.

   b. In the Source field, enter the source IP address, or click the **...** button to choose an IP address that you already defined in ASDM.

      Specify the address and subnet mask using prefix/length notation, such as 10.1.1.0/24. If you enter an IP address without a mask, it is considered to be a host address, even if it ends with a 0.

      Enter **any** to specify any source address.

Separate multiple addresses by a comma.

c. In the Destination field, enter the destination IP address, or click the **...** button to choose an IP address that you already defined in ASDM.

Specify the address and subnet mask using prefix/length notation, such as 10.1.1.0/24. If you enter an IP address without a mask, it is considered to be a host address, even if it ends with a 0.

Enter **any** to specify any destination address.

Separate multiple addresses by a comma.

d. In the Service field, enter an IP service name or number for the destination service, or click the **...** button to choose a service.

If you want to specify a TCP or UDP port number, or an ICMP service number, enter *protocol*/*port*. For example, enter TCP/8080.

By default, the service is IP.

Separate multiple services by a comma.

e. (Optional) Enter a description in the Description field.

f. (Optional) To specify a source service for TCP or UDP, click the **More Options** area open, and enter a TCP or UDP service in the Source Service field.

The destination service and source service must be the same. Copy and paste the destination Service field to the Source Service field.

g. (Optional) To make the rule inactive, click the **More Options** area open, and uncheck **Enable Rule**.

This setting might be useful if you do not want to remove the rule, but want to turn it off.

h. (Optional) To set a time range for the rule, click the **More Options** area open, and from the Time Range drop-down list, choose a time range.

To add a new time range, click the **...** button. See the general operations configuration guide for more information.

This setting might be useful if you only want the rule to be active at predefined times.

- Destination Port—Click **TCP** or **UDP**.

In the Service field, enter a port number or name, or click **...** to choose one already defined in ASDM.

Step 8   Click **Next**.

The Add Management Service Policy Rule - Rule Actions dialog box appears.

Step 9   To configure RADIUS accounting inspection, choose an inspect map from the RADIUS Accounting Map drop-down list, or click **Configure** to add a map.

See Supported Features, page 1-2 for more information.

Step 10   To configure connection settings, see Configuring Connection Settings, page 20-8.

Step 11   Click **Finish**.

Managing the Order of Service Policy Rules

The order of service policy rules on an interface or in the global policy affects how actions are applied to traffic. See the following guidelines for how a packet matches rules in a service policy:

- A packet can match only one rule in a service policy for each feature type.
- When the packet matches a rule that includes actions for a feature type, the ASA does not attempt to match it to any subsequent rules including that feature type.
- If the packet matches a subsequent rule for a different feature type, however, then the ASA also applies the actions for the subsequent rule.

For example, if a packet matches a rule for connection limits, and also matches a rule for application inspection, then both rule actions are applied.

If a packet matches a rule for application inspection, but also matches another rule that includes application inspection, then the second rule actions are not applied.

If your rule includes an ACL with multiple ACEs, then the order of ACEs also affects the packet flow. The ASA tests the packet against each ACE in the order in which the entries are listed. After a match is found, no more ACEs are checked. For example, if you create an ACE at the beginning of an ACL that explicitly permits all traffic, no further statements are ever checked.

To change the order of rules or ACEs within a rule, perform the following steps:

Step 1   From the Configuration > Firewall > Service Policy Rules pane, choose the rule or ACE that you want to move up or down.

Step 2   Click the Move Up or Move Down cursor (see Figure 1-1).

Figure 1-1   Moving an ACE
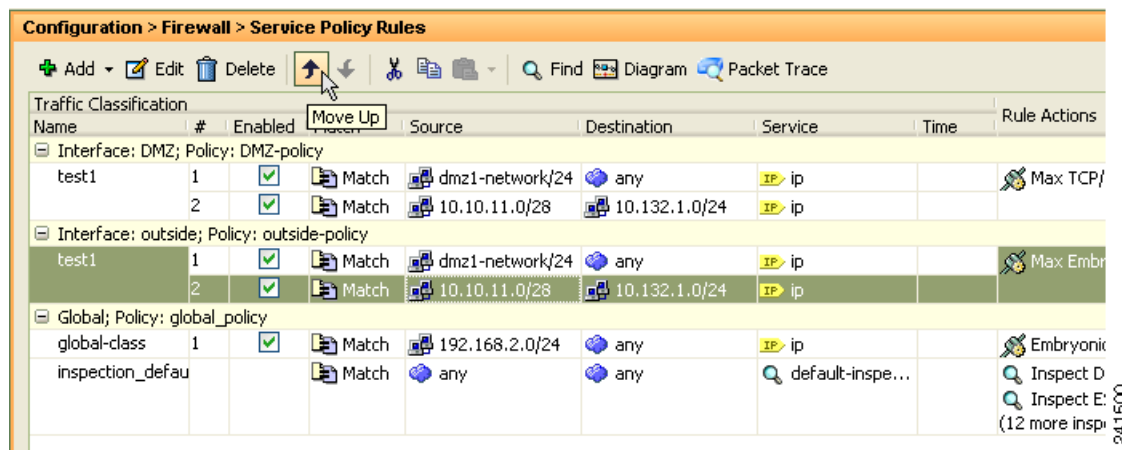
Note   If you rearrange ACEs in an ACL that is used in multiple service policies, then the change is inherited in all service policies.

Step 3   When you are done rearranging your rules or ACEs, click **Apply**.

Feature History for Service Policies

Table 1-3 lists the release history for this feature.

Table 1-3   Feature History for Service Policies

Feature Name	Releases	Feature Information
Modular Policy Framework	7.0(1)	Modular Policy Framework was introduced.
Management class map for use with RADIUS accounting traffic	7.2(1)	The management class map was introduced for use with RADIUS accounting traffic. The following commands were introduced: **class-map type management**, and **inspect radius-accounting**.
Inspection policy maps	7.2(1)	The inspection policy map was introduced. The following command was introduced: **class-map type inspect**.
Regular expressions and policy maps	7.2(1)	Regular expressions and policy maps were introduced to be used under inspection policy maps. The following commands were introduced: **class-map type regex**, **regex**, **match regex**.
Match any for inspection policy maps	8.0(2)	The **match any** keyword was introduced for use with inspection policy maps: traffic can match one or more criteria to match the class map. Formerly, only **match all** was available.

Special Actions for Application Inspections (Inspection Policy Map)

Modular Policy Framework lets you configure special actions for many application inspections. When you enable an inspection engine in the service policy, you can also optionally enable actions as defined in an *inspection policy map*. When the inspection policy map matches traffic within the service policy for which you have defined an inspection action, then that subset of traffic will be acted upon as specified (for example, dropped or rate-limited).

This chapter includes the following sections:

Information About Inspection Policy Maps

See Configuring Application Layer Protocol Inspection, page 8-7 for a list of applications that support inspection policy maps.

An inspection policy map consists of one or more of the following elements. The exact options available for an inspection policy map depends on the application.

- Traffic matching option—You can define a traffic matching option directly in the inspection policy map to match application traffic to criteria specific to the application, such as a URL string, for which you then enable actions.

  - Some traffic matching options can specify regular expressions to match text inside a packet. Be sure to create and test the regular expressions before you configure the policy map, either singly or grouped together in a regular expression class map.

- Inspection class map—An inspection class map includes multiple traffic matching options. You then identify the class map in the policy map and enable actions for the class map as a whole. The difference between creating a class map and defining the traffic match directly in the inspection

policy map is that you can create more complex match criteria and you can reuse class maps. However, you cannot set different actions for different matches. **Note:** Not all inspections support inspection class maps.

- Parameters—Parameters affect the behavior of the inspection engine.

Guidelines and Limitations

- HTTP inspection policy maps—If you modify an in-use HTTP inspection policy map, you must remove and reapply the inspection policy map action for the changes to take effect. For example, if you modify the "http-map" inspection policy map, you must remove, apply changes, and readd the inspection policy map to the service policy.

- All inspection policy maps—If you want to exchange an in-use inspection policy map for a different map name, you must remove, apply changes, and readd the new inspection policy map to the service policy.

- You can specify multiple inspection class maps or direct matches in the inspection policy map.

  If a packet matches multiple different matches, then the order in which the ASA applies the actions is determined by internal ASA rules, and not by the order they are added to the inspection policy map. The internal rules are determined by the application type and the logical progression of parsing a packet, and are not user-configurable. For example for HTTP traffic, parsing a Request Method field precedes parsing the Header Host Length field; an action for the Request Method field occurs before the action for the Header Host Length field.

  If an action drops a packet, then no further actions are performed in the inspection policy map. For example, if the first action is to reset the connection, then it will never match any further match criteria. If the first action is to log the packet, then a second action, such as resetting the connection, can occur.

  If a packet matches multiple match criteria that are the same, then they are matched in the order they appear in the policy map.

  A class map is determined to be the same type as another class map or direct match based on the lowest priority match option in the class map (the priority is based on the internal rules). If a class map has the same type of lowest priority match option as another class map, then the class maps are matched according to the order they are added to the policy map. If the lowest priority match for each class map is different, then the class map with the higher priority match option is matched first.

Default Inspection Policy Maps

DNS inspection is enabled by default, using the preset_dns_map inspection class map:

- The maximum DNS message length is 512 bytes.

- The maximum client DNS message length is automatically set to match the Resource Record.

- DNS Guard is enabled, so the ASA tears down the DNS session associated with a DNS query as soon as the DNS reply is forwarded by the ASA. The ASA also monitors the message exchange to ensure that the ID of the DNS reply matches the ID of the DNS query.

- Translation of the DNS record based on the NAT configuration is enabled.

- Protocol enforcement is enabled, which enables DNS message format check, including domain name length of no more than 255 characters, label length of 63 characters, compression, and looped pointer check.

> **Note**   There are other default inspection policy maps such as **_default_esmtp_map**. For example, an ESMTP inspection rule implicitly uses the policy map "_default_esmtp_map."

Defining Actions in an Inspection Policy Map

When you enable an inspection engine in the service policy, you can also optionally enable actions as defined in an inspection policy map.

Detailed Steps

Step 1   (Optional) Create an inspection class map. Alternatively, you can identify the traffic directly within the policy map. See Identifying Traffic in an Inspection Class Map, page 2-3.

Step 2   (Optional) For policy map types that support regular expressions, create a regular expression. See the general operations configuration guide.

Step 3   Choose **Configuration > Firewall > Objects > Inspect Maps** .

Step 4   Choose the inspection type you want to configure.

Step 5   Click **Add** to add a new inspection policy map.

Step 6   Follow the instructions for your inspection type in the inspection chapter.

Identifying Traffic in an Inspection Class Map

This type of class map allows you to match criteria that is specific to an application. For example, for DNS traffic, you can match the domain name in a DNS query.

A class map groups multiple traffic matches (in a match-all class map), or lets you match any of a list of matches (in a match-any class map). The difference between creating a class map and defining the traffic match directly in the inspection policy map is that the class map lets you group multiple match commands, and you can reuse class maps. For the traffic that you identify in this class map, you can specify actions such as dropping, resetting, and/or logging the connection in the inspection policy map. If you want to perform different actions on different types of traffic, you should identify the traffic directly in the policy map.

Restrictions

Not all applications support inspection class maps.

Detailed Steps

Step 1   Choose **Configuration > Firewall > Objects > Class Maps** .

Step 2   Choose the inspection type you want to configure.

Step 3   Click **Add** to add a new inspection class map.

Step 4   Follow the instructions for your inspection type in the inspection chapter.

Where to Go Next

To use an inspection policy, see Chapter 1, "Service Policy."

Feature History for Inspection Policy Maps

Table 2-1 lists the release history for this feature.

Table 2-1   Feature History for Service Policies

Feature Name	Releases	Feature Information
Inspection policy maps	7.2(1)	The inspection policy map was introduced. The following command was introduced: **class-map type inspect**.
Regular expressions and policy maps	7.2(1)	Regular expressions and policy maps were introduced to be used under inspection policy maps. The following commands were introduced: **class-map type regex**, **regex**, **match regex**.
Match any for inspection policy maps	8.0(2)	The **match any** keyword was introduced for use with inspection policy maps: traffic can match one or more criteria to match the class map. Formerly, only **match all** was available.

Access Rules

This chapter describes how to control network access through the ASA using access rules and includes the following sections:

- Information About Access Rules, page 3-1
- Licensing Requirements for Access Rules, page 3-6
- Guidelines and Limitations, page 3-7
- Default Settings, page 3-7
- Configuring Access Rules, page 3-7
- Feature History for Access Rules, page 3-13

Note   You use access rules to control network access in both routed and transparent firewall modes. In transparent mode, you can use both access rules (for Layer 3 traffic) and EtherType rules (for Layer 2 traffic).

To access the ASA interface for management access, you do not also need an access rule allowing the host IP address. You only need to configure management access according to the general operations configuration guide.

Information About Access Rules

Your access policy is made up of one or more access rules and/or EtherType rules per interface or globally for all interfaces.

You can use access rules in routed and transparent firewall mode to control IP traffic. An access rule permits or denies traffic based on the protocol, a source and destination IP address or network, and optionally the source and destination ports.

For transparent mode only, an EtherType rule controls network access for non-IP traffic. An EtherType rule permits or denies traffic based on the EtherType.

This section includes the following topics:

- General Information About Rules, page 3-2
- Information About Access Rules, page 3-4
- Information About EtherType Rules, page 3-5

General Information About Rules

This section describes information for both access rules and EtherType rules, and it includes the following topics:

Implicit Permits

For routed mode, the following types of traffic are allowed through by default:

- Unicast IPv4 traffic from a higher security interface to a lower security interface.
- Unicast IPv6 traffic from a higher security interface to a lower security interface.

For transparent mode, the following types of traffic are allowed through by default:

- Unicast IPv4 traffic from a higher security interface to a lower security interface.
- Unicast IPv6 traffic from a higher security interface to a lower security interface.
- ARPs in both directions.

> **Note** ARP traffic can be controlled by ARP inspection, but cannot be controlled by an access rule.

- BPDUs in both directions.

For other traffic, you need to use either an access rule (IPv4 and IPv6) or an EtherType rule (non-IPv4/IPv6).

Information About Interface Access Rules and Global Access Rules

You can apply an access rule to a specific interface, or you can apply an access rule globally to all interfaces. You can configure global access rules in conjunction with interface access rules, in which case, the specific interface access rules are always processed before the general global access rules.

> **Note** Global access rules apply only to inbound traffic. See Inbound and Outbound Rules, page 3-3.

Using Access Rules and EtherType Rules on the Same Interface

You can apply both access rules and EtherType rules to each direction of an interface.

Rule Order

The order of rules is important. When the ASA decides whether to forward or drop a packet, the ASA tests the packet against each rule in the order in which the rules are listed. After a match is found, no more rules are checked. For example, if you create an access rule at the beginning that explicitly permits all traffic for an interface, no further rules are ever checked. For more information, see Implicit Deny, page 3-3.

You can disable a rule by making it inactive.

Implicit Deny

ACLs have an implicit deny at the end of the list, so unless you explicitly permit it, traffic cannot pass. For example, if you want to allow all users to access a network through the ASA except for particular addresses, then you need to deny the particular addresses and then permit all others.

For EtherType ACLs, the implicit deny at the end of the ACL does not affect IP traffic or ARPs; for example, if you allow EtherType 8037, the implicit deny at the end of the ACL does not now block any IP traffic that you previously allowed with an extended ACL (or implicitly allowed from a high security interface to a low security interface). However, if you explicitly deny all traffic with an EtherType ACE, then IP and ARP traffic is denied.

If you configure a global access rule, then the implicit deny comes *after* the global rule is processed. See the following order of operations:

1. Interface access rule.
2. Global access rule.
3. Implicit deny.

Using Remarks

In the ASDM access rule window, a remark that displays next to the rule is the one that was configured before the rule, so when you configure a remark from the CLI and then view it in an ASDM access rule window, the remark displays next to the rule that was configured after the remark in the CLI. However, the packet tracer in ASDM matches the remark that is configured after the matching rule in the CLI.

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).

Inbound and Outbound Rules

The ASA supports two types of ACLs:

• Inbound—Inbound access rules apply to traffic as it enters an interface. Global access rules are always inbound.

• Outbound—Outbound ACLs apply to traffic as it exits an interface.

> **Note** "Inbound" and "outbound" refer to the application of an ACL on an interface, either to traffic entering the ASA on an interface or traffic exiting the ASA on an interface. These terms do not refer to the movement of traffic from a lower security interface to a higher security interface, commonly known as inbound, or from a higher to lower interface, commonly known as outbound.

An outbound ACL is useful, for example, if you want to allow only certain hosts on the inside networks to access a web server on the outside network. Rather than creating multiple inbound ACLs to restrict access, you can create a single outbound ACL that allows only the specified hosts. (See Figure 3-1.) The outbound ACL prevents any other hosts from reaching the outside network.

Figure 3-1   Outbound ACL
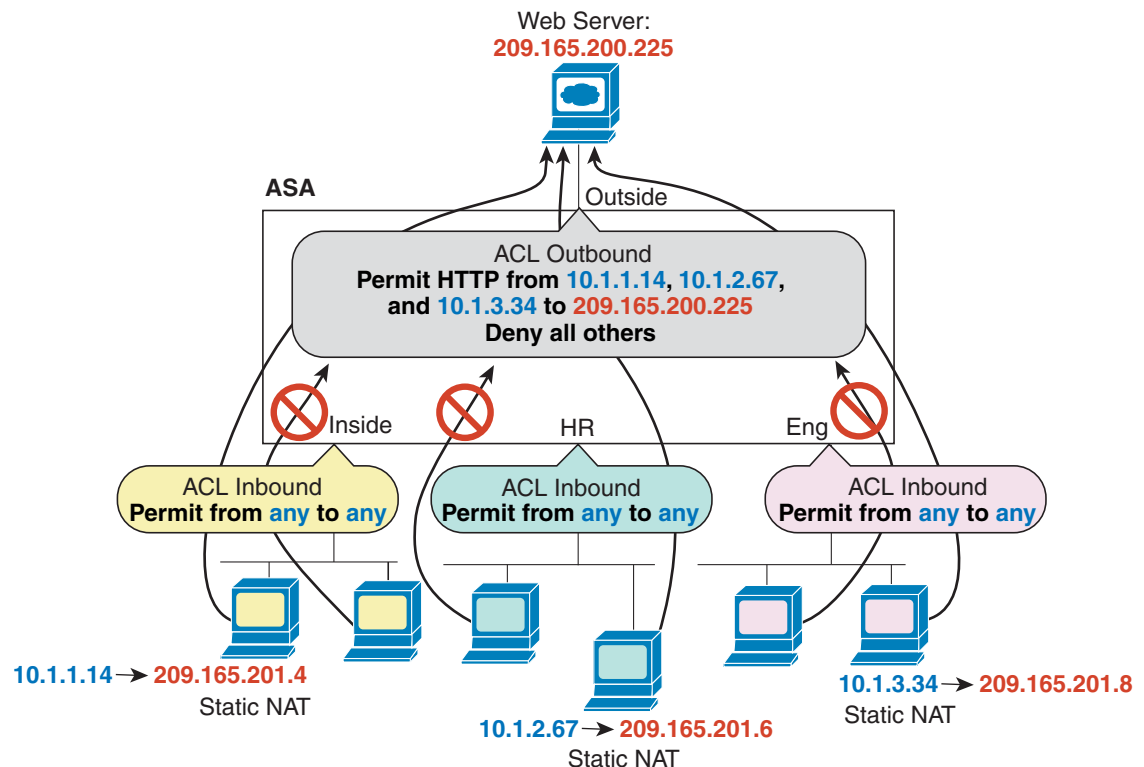
Information About Access Rules

This section describes information about access rules and includes the following topics:

- Access Rules for Returning Traffic, page 3-5
- Allowing Broadcast and Multicast Traffic through the Transparent Firewall Using Access Rules, page 3-5
- Management Access Rules, page 3-5

Access Rules for Returning Traffic

For TCP and UDP connections for both routed and transparent mode, you do not need an access rule to allow returning traffic because the ASA allows all returning traffic for established, bidirectional connections.

For connectionless protocols such as ICMP, however, the ASA establishes unidirectional sessions, so you either need access rules to allow ICMP in both directions (by applying ACLs to the source and destination interfaces), or you need to enable the ICMP inspection engine. The ICMP inspection engine treats ICMP sessions as bidirectional connections.

Allowing Broadcast and Multicast Traffic through the Transparent Firewall Using Access Rules

In routed firewall mode, broadcast and multicast traffic is blocked even if you allow it in an access rule, including unsupported dynamic routing protocols and DHCP (unless you configure DHCP relay). Transparent firewall mode can allow any IP traffic through.

Note Because these special types of traffic are connectionless, you need to apply an access rule to both interfaces, so returning traffic is allowed through.

Table 3-1 lists common traffic types that you can allow through the transparent firewall.

Table 3-1   *Transparent Firewall Special Traffic*

Traffic Type	Protocol or Port	Notes
DHCP	UDP ports 67 and 68	If you enable the DHCP server, then the ASA does not pass DHCP packets.
EIGRP	Protocol 88	—
OSPF	Protocol 89	—
Multicast streams	The UDP ports vary depending on the application.	Multicast streams are always destined to a Class D address (224.0.0.0 to 239.x.x.x).
RIP (v1 or v2)	UDP port 520	—

Management Access Rules

You can configure access rules that control management traffic destined to the ASA. Access control rules for to-the-box management traffic (such as HTTP, Telnet, and SSH) have higher precedence than an management access rule. Therefore, such permitted management traffic will be allowed to come in even if explicitly denied by the to-the-box ACL.

Information About EtherType Rules

This section describes EtherType rules and includes the following topics:

- Supported EtherTypes and Other Traffic, page 3-6
- Access Rules for Returning Traffic, page 3-6
- Allowing MPLS, page 3-6

Supported EtherTypes and Other Traffic

An EtherType rule controls the following:

- EtherType identified by a 16-bit hexadecimal number, including common types IPX and MPLS unicast or multicast.
- Ethernet V2 frames.
- BPDUs, which are permitted by default. BPDUs are SNAP-encapsulated, and the ASA is designed to specifically handle BPDUs.
- Trunk port (Cisco proprietary) BPDUs. Trunk BPDUs have VLAN information inside the payload, so the ASA modifies the payload with the outgoing VLAN if you allow BPDUs.
- IS-IS.

The following types of traffic are not supported:

- 802.3-formatted frames—These frames are not handled by the rule because they use a length field as opposed to a type field.

Access Rules for Returning Traffic

Because EtherTypes are connectionless, you need to apply the rule to both interfaces if you want traffic to pass in both directions.

Allowing MPLS

If you allow MPLS, ensure that Label Distribution Protocol and Tag Distribution Protocol TCP connections are established through the ASA by configuring both MPLS routers connected to the ASA to use the IP address on the ASA interface as the router-id for LDP or TDP sessions. (LDP and TDP allow MPLS routers to negotiate the labels (addresses) used to forward packets.)

On Cisco IOS routers, enter the appropriate command for your protocol, LDP or TDP. The *interface* is the interface connected to the ASA.

```
hostname(config)# mpls ldp router-id interface force
```

Or

```
hostname(config)# tag-switching tdp router-id interface force
```

Licensing Requirements for Access Rules

Model	License Requirement
ASAv	Standard or Premium License.
All other models	Base License.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall modes.

IPv6 Guidelines

Supports IPv6. (9.0 and later) The source and destination addresses can include any mix of IPv4 and IPv6 addresses. For pre-9.0 versions, you must create a separate IPv6 access rule.

Additional Guidelines and Limitations

- You can reduce the memory required to search access rules by enabling object group search, but this is at the expense rule lookup performance. When enabled, object group search does not expand network objects, but instead searches access rules for matches based on those group definitions. You can set this option by clicking the **Advanced** button below the access rule table.

- You can improve system performance and reliability by using the transactional commit model for access groups. See the basic settings chapter in the general operations configuration guide for more information. The option is under **Configurations > Device Management > Advanced > Rule Engine**.

Default Settings

See Implicit Permits, page 3-2.

Configuring Access Rules

This section includes the following topics:

Adding an Access Rule

To apply an access rule, perform the following steps.

Detailed Steps

Step 1 Choose **Configuration** > **Firewall** > **Access Rules**.

Step 2 Click **Add**, and choose one of the following options:

The Add Access Rule dialog box appears.

Step 3 From the Interface drop-down list, choose the interface on which to apply the rule. Choose **Any** to apply a global rule.

Step 4 In the Action field, click one of the following radio buttons next to the desired action:

- **Permit**—Permits access if the conditions are matched.
- **Deny**—Denies access if the conditions are matched.

Step 5 In the Source field, enter an IP address that specifies the network, interface IP, or any address from which traffic is permitted or denied to the specified destination. You may use either an IPv4 or IPv6 address.

For more information about enabling IPv6 on an interface, see the general operations configuration guide.

Step 6 In the User field, enter a user name or group to the ACL. Enter the user name in the format *domain_NetBIOS_name\user_name.* Enter the group name in the format *domain_NetBIOS_name\group_name.*

You can configure access rules based on user names and user group names rather than through source IP addresses. The ASA applies the security policies based on an association of IP addresses to Windows Active Directory login information and reports events based on the mapped user names instead of network IP addresses.

See the general operations configuration guide for more information.

Step 7 To browse for a user name or user group, click the ellipsis (...) button. The Browse User dialog box appears.

Step 8 In the Destination field, enter an IP address that specifies the network, interface IP, any address to which traffic is permitted or denied from the source specified in the Source field. You may use either an IPv4 or IPv6 address.

Step 9 Select the service type.

Step 10 (Optional) To add a time range to your access rule that specifies when traffic can be allowed or denied, click **More Options** to expand the list.

    a. To the right of the Time Range drop down list, click the browse button.

       The Browse Time Range dialog box appears.

    b. Click **Add**.

       The Add Time Range dialog box appears.

    c. In the Time Range Name field, enter a time range name, with no spaces.

    d. Choose the Start Time and the End Time.

    e. To specify additional time constraints for the time range, such as specifying the days of the week or the recurring weekly interval in which the time range will be active, click **Add**, and choose the specifications.

    f. Click **OK** to apply the optional time range specifications.

Step 11 (Optional) In the Description field, add a text description about the access rule.

The description can contain multiple lines; however, each line can be no more than 100 characters in length.

Step 12 (Optional) Logging is enabled by default. You can disable logging by unchecking the check box, or you can change the logging level from the drop-down list. The default logging level is Informational.

Step 13 Click **OK**. The access rule appears with the newly configured access rules.

Step 14 Click **Apply** to save the access rule to your configuration.

You can edit or delete a particular access rule by selecting the rule and then clicking Edit or Delete.

Adding an EtherType Rule (Transparent Mode Only)

The EtherType Rules window shows access rules based on packet EtherTypes. EtherType rules are used to configure non-IP related traffic policies through the ASA when operating in transparent mode. In transparent mode, you can apply both extended and EtherType access rules to an interface. EtherType rules take precedence over the extended access rules.

For more information about EtherType rules, see Information About Access Rules, page 3-1.

To add an EtherType rule, perform the following steps:

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **EtherType Rules**.

Step 2 Click **Add**.

The Add EtherType rules window appears.

Step 3 (Optional) To specify the placement of the new EtherType rule, select an existing rule, and click **Insert...** to add the EtherType rule before the selected rule, or click **Insert After...** to add the EtherType rule after the selected rule.

Step 4 From the Interface drop-down list, choose the interface on which to apply the rule. Choose **Any** to apply a global rule.

Step 5 In the Action field, click one of the following radio buttons next to the desired action:

- **Permit**—Permits access if the conditions are matched.
- **Deny**—Denies access if the conditions are matched.

Step 6 In the EtherType field, choose an EtherType value from the drop-down list.

Step 7 (Optional) In the Description field, add a test description about the rule.

The description can contain multiple lines; however, each line can be no more than 100 characters in length.

Step 8 (Optional) To specify the direction for this rule, click **More Options** to expand the list, and then specify the direction by clicking one of the following radio buttons:

- **In**—Incoming traffic
- **Out**—Outgoing traffic

Step 9 Click **OK**.

Configuring Management Access Rules

You can configure an interface ACL that supports access control for to-the-box management traffic from a specific peer (or set of peers) to the security appliance. One scenario in which this type of ACL would be useful is when you want to block IKE Denial of Service attacks.

To configure an extended ACL that permits or denies packets for to-the-box traffic, perform the following steps:

Step 1   Choose **Configuration** > **Device Management** > **Management Access** > **Management Access Rules**.

Step 2   Click **Add**, and choose one of the following actions:

The Add Management Access Rule dialog box appears.

Step 3   From the Interface drop-down list, choose an interface on which to apply the rule. Choose **Any** to apply a global rule.

Step 4   In the Action field, click one of the following radio buttons to choose the action:

- **Permit**—Permits access if the conditions are matched.
- **Deny**—Denies access if the conditions are matched.

Step 5   In the Source field, enter an IP address that specifies the network object group, interface IP, or any address from which traffic is permitted or denied. You may use either an IPv4 or IPv6 address.

> ✎
>
> **Note**   IPv6 must be enabled on at least one interface before you can configure an extended ACL with an IPv6 address. For more information about enabling IPv6 on an interface, see the general operations configuration guide.

Step 6   In the Service field, add a service name for rule traffic, or click the ellipsis (...) to browse for a service.

Step 7   (Optional) In the Description field, add a description for this management access rule.

The description can contain multiple lines; however, each line can be no more than 100 characters in length.

Step 8   (Optional) Logging is enabled by default. You can disable logging by unchecking the check box, or you can change the logging level from the drop-down list. The default logging level is Informational.

Step 9   (Optional) To add a source service (TCP, UDP, and TCP-UDP only) and a time range to your access rule that specifies when traffic can be allowed or denied, click **More Options** to expand the list.If you want to turn off this Management Access Rule, uncheck **Enable Rule**.

- Add a source service in the Source Service field, or click the ellipsis (...) to browse for a service.

  The destination service and source service must be the same. Copy and paste the destination Service field to the Source Service field.

- To configure the logging interval (if you enable logging and choose a non-default setting), enter a value in seconds in the Logging Interval field.

- To select a predefined time range for this rule, from the Time Range drop-down list, choose a time range; or click the ellipsis (...) to browse for a time range. You can also specify additional time constraints for the time range, such as specifying the days of the week or the recurring weekly interval in which the time range will be active.

Step 10   Click **OK**. The dialog box closes, and the Management Access rule is added.

Step 11   Click **Apply**. The rule is saved in the running configuration.

Advanced Access Rule Configuration

The Advanced Access Rule Configuration dialog box lets you to set global access rule logging options.

When you enable logging, if a packet matches the access rule, the ASA creates a flow entry to track the number of packets received within a specific interval. The ASA generates a system log message at the first hit and at the end of each interval, identifying the total number of hits during the interval and reporting the time of the last hit.

> **Note**   The ASApane displays the hit count information in the "last rule hit" row. To view the rule hit count and timestamp, choose **Configuration** > **Firewall** > **Advanced** > **ACL Manager**, and hover the mouse pointer over a cell in the ACL Manager table.

At the end of each interval, the ASA resets the hit count to 0. If no packets match the access rule during an interval, the ASA deletes the flow entry.

A large number of flows can exist concurrently at any point of time. To prevent unlimited consumption of memory and CPU resources, the ASA places a limit on the number of concurrent deny flows; the limit is placed only on deny flows (and not permit flows) because they can indicate an attack. When the limit is reached, the ASA does not create a new deny flow until the existing flows expire. If someone initiates a denial of service attack, the ASA can create a very large number of deny flows in a very short period of time. Restricting the number of deny-flows prevents unlimited consumption of memory and CPU resources.

Prerequisites

These settings only apply if you enable the newer logging mechanism for the access rule.

Fields

- Maximum Deny-flows—The maximum number of deny flows permitted before the ASA stops logging, between 1 and the default value. The default is 4096.

- Alert Interval—The amount of time (1-3600 seconds) between system log messages (number 106101) that identify that the maximum number of deny flows was reached. The default is 300 seconds.

- Per User Override table—Specifies the state of the per user override feature. If the per user override feature is enabled on the inbound access rule, the access rule provided by a RADIUS server replaces the access rule configured on that interface. If the per user override feature is disabled, the access rule provided by the RADIUS server is combined with the access rule configured on that interface. If the inbound access rule is not configured for the interface, per user override cannot be configured.

  By default, VPN remote access traffic is not matched against interface ACLs. However, if you deselect the **Enable inbound VPN sessions to bypass interface access lists** setting on the Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Connection Profiles pane), the behavior depends on whether there is a VPN filter applied in the group policy (see the Configuration > Remote Access VPN > Network (Client) Access > Group Policies > Add/Edit > General > More Options > Filter field) and whether you set the Per User Override option:

  – No Per User Override, no VPN filter —Traffic is matched against the interface ACL.

>   – No Per User Override, VPN filter —Traffic is matched first against the interface ACL, then against the VPN filter.

>   – Per User Override, VPN filter —Traffic is matched against the VPN filter only.

- Object Group Search Setting—Reduces the amount of memory used to store service rules, but lengthens the amount of time to search for a matching access rule.

Access Rule Explosion

The security appliance allows you to turn off the expansion of access rules that contain certain object groups. When expansion is turned off, an object group search is used for lookup, which lowers the memory requirements for storing expanded rules but decreases the lookup performance. Because of the trade-off of performance for memory utilization, you can turn on and turn off the search.

To configure the option of turning off the expansion of access rules that contain s, perform the following steps:

Step 1   Choose **Configuration** > **Firewall** > **Access Rules**.

Step 2   Click the **Advanced** button.

Step 3   Check the **Enable Object Group Search Algorithm** check box.

Configuring HTTP Redirect

The HTTP Redirect table displays each interface on the ASA, shows whether it is configured to redirect HTTP connections to HTTPS, and the port number from which it redirects those connections.

Note   To redirect HTTP, the interface requires an ACL that permits HTTP. Otherwise, the interface cannot listen to the HTTP port.

The Configuration > Device Management > Advanced > HTTP Redirect > Edit pane lets you change the HTTP redirect setting of an interface or the port from which it redirects HTTP connections. Select the interface in the table and click **Edit**. You can also double-click an interface. The Edit HTTP/HTTPS Settings dialog box opens.

Edit HTTP/HTTPS Settings

The Edit HTTP/HTTPS Settings dialog box lets you change the HTTP redirect setting of an interface or the port number.

Fields

The Edit HTTP/HTTPS Settings dialog box includes the following fields:

- Interface—Identifies the interface on which the ASA redirects or does not redirect HTTP requests to HTTPS.

- Redirect HTTP to HTTPS—Check to redirect HTTP requests to HTTPS, or uncheck to not redirect HTTP requests to HTTPS.

- **HTTP Port**—Identifies the port from which the interface redirects HTTP connections. By default it listens to port 80.

For more information about access rules, see Information About Access Rules, page 3-1.

Feature History for Access Rules

Table 3-2 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 3-2   Feature History for Access Rules

Feature Name	Platform Releases	Feature Information
Interface access rules	7.0(1)	Controlling network access through the ASA using ACLs. We introduced the following screen: Configuration > Firewall > Access Rules.
Global access rules	8.3(1)	Global access rules were introduced. We modified the following screen: Configuration > Firewall > Access Rules.
Support for Identity Firewall	8.4(2)	You can now use identity firewall users and groups for the source and destination. You can use an identity firewall ACL with access rules, AAA rules, and for VPN authentication.
EtherType ACL support for IS-IS traffic	8.4(5), 9.1(2)	In transparent firewall mode, the ASA can now pass IS-IS traffic using an EtherType ACL. We modified the following screen: Configuration > Device Management > Management Access > EtherType Rules.
Support for TrustSec	9.0(1)	You can now use TrustSec security groups for the source and destination. You can use an identity firewall ACL with access rules.

Table 3-2   *Feature History for Access Rules (continued)*

Feature Name	Platform Releases	Feature Information
Unified ACL for IPv4 and IPv6	9.0(1)	ACLs now support IPv4 and IPv6 addresses. You can even specify a mix of IPv4 and IPv6 addresses for the source and destination. The **any** keyword was changed to represent IPv4 and IPv6 traffic. The **any4** and **any6** keywords were added to represent IPv4-only and IPv6-only traffic, respectively. The IPv6-specific ACLs are deprecated. Existing IPv6 ACLs are migrated to extended ACLs. See the release notes for more information about migration. We modified the following screens: Configuration > Firewall > Access Rules Configuration > Remote Access VPN > Network (Client) Access > Group Policies > General > More Options
Extended ACLand object enhancement to filter ICMP traffic by ICMP code	9.0(1)	ICMP traffic can now be permitted/denied based on ICMP code. We introduced or modified the following screens: Configuration > Firewall > Objects > Service Objects/Groups Configuration > Firewall > Access Rule

Public Servers

This section describes how to configure public servers, and includes the following topics:

Information About Public Servers

The Public Servers pane enables an administrator to provide internal and external users access to various application servers. This pane displays a list of public servers. internal and external addresses, the interfaces to which the internal or external addresses apply, the ability to translate the addresses, and the service that is exposed. You can add, edit, delete, or modify settings for existing public servers.

Licensing Requirements for Public Servers

Model	License Requirement
ASAv	Standard or Premium License.
All other models	Base License.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Adding a Public Server that Enables Static NAT

To add a public server that enables static NAT and creates a fixed translation of a real address to a mapped address, perform the following steps:

Step 1   In the Configuration > Firewall > Public Servers pane, click **Add** to add a new server.

The Add Public Server dialog box appears.

Step 2   From the Private Interface drop-down menu, select the name of the private interface to which the real server is connected.

Step 3   In the Private IP address field, enter the real IP address of the server (IPv4 only).

Step 4   In the Private Service field, click **Browse** to display the Browse Service dialog box, choose the actual service that is exposed to the outside, and click **OK**.

Optionally, from the Browse Service dialog box you can click **Add** to create a new service or service group. Multiple services from various ports can be opened to the outside. For more information about service objects and service groups, see the general operations configuration guide.

Step 5   From the Public Interface drop-down menu, enter the interface through which users from the outside can access the real server.

Step 6   In the Public Address field, enter the mapped IP address of the server, which is the address that is seen by the outside user.

Step 7   (Optional) To enable static PAT, check the **Specify if Public Service is different from private service** check box .

Step 8   Click **OK**. The configuration appears in the main pane.

Step 9   Click **Apply** to generate static NAT and a corresponding access rule for the traffic flow and to save the configuration.

For information about static NAT, see .

Adding a Public Server that Enables Static NAT with PAT

To add a public server that lets you specify a real and mapped protocol (TCP or UDP) to a port, perform the following steps:

Step 1   Choose **Configuration > Firewall > Public Servers**, then click **Add**.

The Add Public Server dialog box appears.

Step 2 From the Private Interface drop-down menu, choose the name of the private interface to which the real server is connected.

Step 3 In the Private IP address field, enter the real IP address of the server (only IPv4 is supported).

Step 4 In the Private Service field, click **Browse** to display the Browse Service dialog box

Step 5 Choose the actual service that is exposed to the outside, and click **OK**.

Optionally, from the Browse Service dialog box, click **Add** to create a new service or service group. Multiple services from various ports can be opened to the outside. For more information about service objects and service groups, see the general operations configuration guide.

Step 6 From the Public Interface drop-down menu, enter the interface through which users from the outside can access the real server.

Step 7 In the Public Address field, enter the mapped IP address of the server, which is the address that the outside user sees.

Step 8 Check the **Specify Public Service if different from Private Service** check box to enable static PAT.

Step 9 In the Public Service field, enter the mapped protocol (TCP or UDP only), or click **Browse** to choose a protocol from the list.

Step 10 Click **OK**.

Step 11 Click **Apply** to generate static NAT with PAT and a corresponding access rule for the traffic flow, and to save the configuration.

For information about static NAT with port address translation, see Information About Static NAT with Port Translation, page 5-4.

Editing Settings for a Public Server

To edit the settings for a public server, perform the following steps:

Step 1 Choose **Configuration > Firewall > Public Servers**, choose an existing public server, then click **Edit**.

The Edit Public Server dialog box appears.

Step 2 Make any necessary changes to the following settings:

- Private Interface—The interface to which the real server is connected.
- Private IP Address—The real IP address of the server.
- Private Service—The actual service that is running on the real server.
- Public Interface—The interface through which outside users can access the real server.
- Public Address—The IP address that is seen by outside users.
- Public Service—The service that is running on the translated address. Click the **Information** icon to view information about supported public services.

Step 3 Click **OK**, then click **Apply** to save your changes.

Feature History for Public Servers

Table 4-1 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 4-1   Feature History for Public Servers

Feature Name	Platform Releases	Feature Information
Public Servers	8.3(1)	Public servers provide internal and external users access to various application servers. We introduced the following screen: Configuration > Firewall > Public Servers

P A R T 2

Network Address Translation

Information About NAT (ASA 8.3 and Later)

This chapter provides an overview of how Network Address Translation (NAT) works on the ASA. This chapter includes the following sections:

Note To start configuring NAT, see Chapter 6, "Network Object NAT (ASA 8.3 and Later)," or Chapter 7, "Twice NAT (ASA 8.3 and Later)."

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) (Version 9.0(1) and later)—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 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 ASA, 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 residing on the ASA's 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

NAT Types Overview

You can implement NAT using the following methods:

- Static NAT—A consistent mapping between a real and mapped IP address. Allows bidirectional traffic initiation. See Static NAT, page 5-3.

- 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, page 5-8.

- 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, page 5-10.

- 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, page 5-12.

Static NAT

This section describes static NAT and includes the following topics:

Information 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.

Figure 5-1 shows a typical static NAT scenario. The translation is always active so both real and remote hosts can initiate connections.

Figure 5-1   ***Static NAT***

Security
Appliance

10.1.1.1   →   209.165.201.1

10.1.1.2   →   209.165.201.2

Inside / Outside

130035

Note   You can disable bidirectionality if desired.

Information About Static NAT with Port Translation

Static NAT with port translation lets you specify a real and mapped protocol (TCP or UDP) and port.

This section includes the following topics:

- Information About Static NAT with Port Address Translation, page 5-4
- Static NAT with Identity Port Translation, page 5-5
- Static NAT with Port Translation for Non-Standard Ports, page 5-5
- Static Interface NAT with Port Translation, page 5-6

Information About Static NAT with Port Address Translation

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.

Figure 5-2 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 5-2   ***Typical Static NAT with Port Translation Scenario***

Security
Appliance

10.1.1.1:23   →   209.165.201.1:23

10.1.1.2:8080   →   209.165.201.2:80

Inside / Outside

130044

> **Note**   For applications that require application inspection for secondary channels (for example, FTP and VoIP), the ASA automatically translates the secondary ports.

Static NAT with Identity 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. (See Figure 5-3. See Single Address for FTP, HTTP, and SMTP (Static NAT-with-Port-Translation), page 6-32 for details on how to configure this example.)

Figure 5-3   *Static NAT with Port Translation*
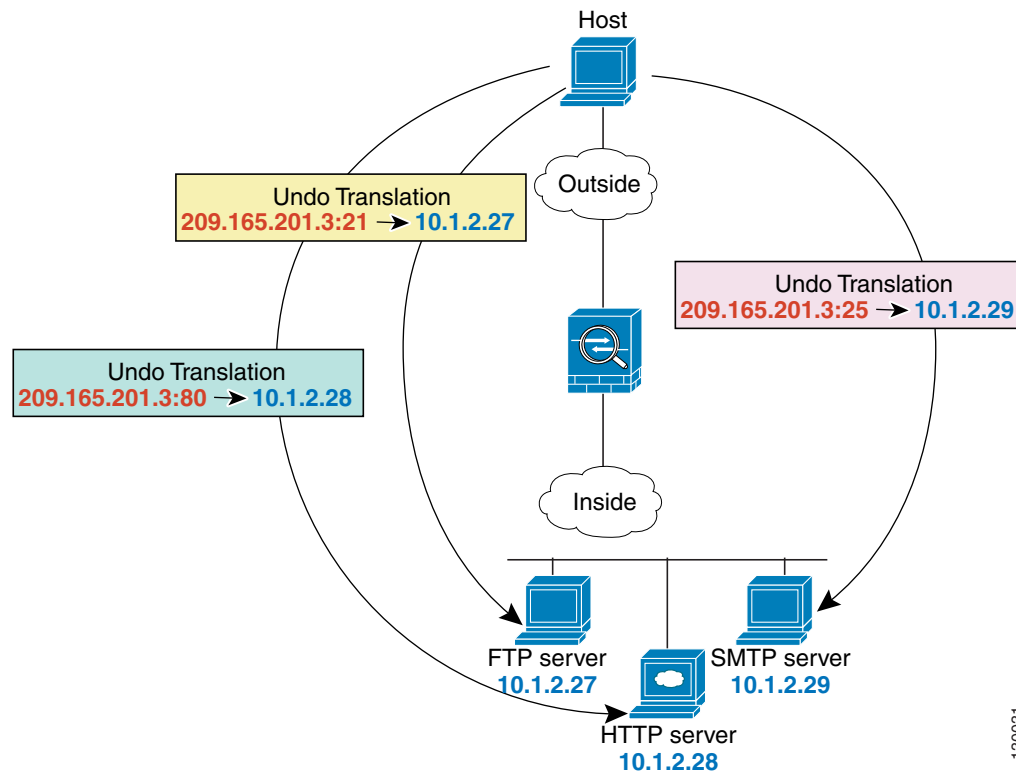
Static NAT with Port Translation for Non-Standard Ports

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 ASA outside interface to an inside host, then you can map the inside host IP address/port 23 to the ASA interface address/port 23. (Note that although Telnet to the ASA is not allowed to the lowest security interface, static NAT with interface port translation redirects the Telnet session instead of denying it).

Information About 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.

Figure 5-4 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/1st mapped IP is technically the only bidirectional translation.

Figure 5-4   *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 (see Figure 5-5). (See Inside Load Balancer with Multiple Mapped Addresses (Static NAT, One-to-Many), page 6-28 for details on how to configure this example.)

Figure 5-5   *One-to-Many Static NAT*
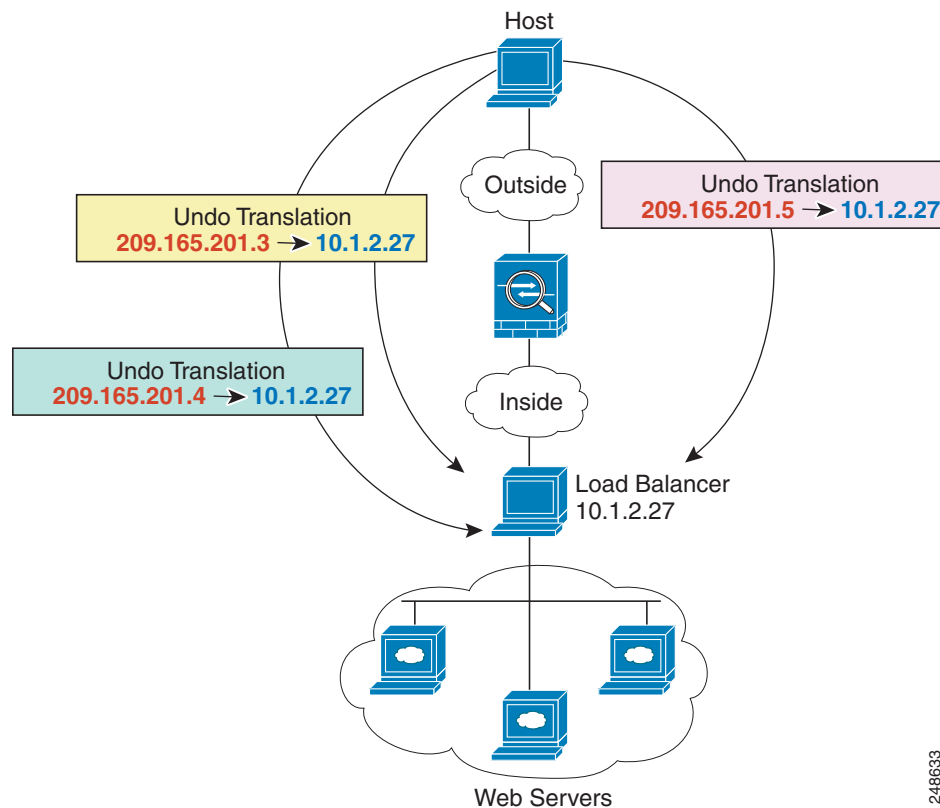
Information About Other Mapping Scenarios (Not Recommended)

The ASA 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.

Figure 5-6 shows a typical few-to-many static NAT scenario.

Figure 5-6   **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).

Figure 5-7 shows a typical many-to-few static NAT scenario.

Figure 5-7   **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.

Dynamic NAT

This section describes dynamic NAT and includes the following topics:

- Information About Dynamic NAT, page 5-9
- Dynamic NAT Disadvantages and Advantages, page 5-10

Information 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, the ASA 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.

Figure 5-8 shows a typical dynamic NAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back.

Figure 5-8   *Dynamic NAT*
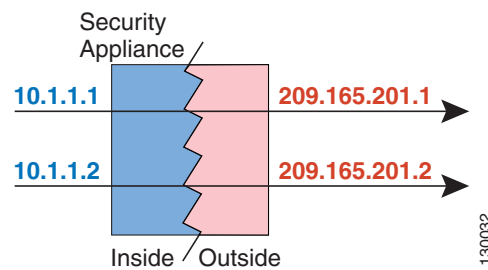
Figure 5-9 shows a remote host attempting to initiate a connection to a mapped address. This address is not currently in the translation table; therefore, the ASA drops the packet.

Figure 5-9   *Remote Host Attempts to Initiate a Connection to a Mapped Address*
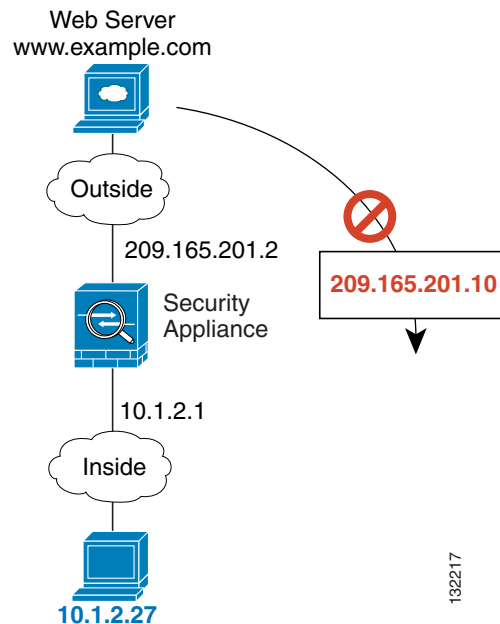
> **Note** For the duration of the translation, a remote host can initiate a connection to the translated host if an access rule allows it. Because the address is unpredictable, a connection to the host is unlikely. Nevertheless, in this case you can rely on the security of the access rule.

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 fallback 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.

See Default Settings and NAT Limitations, page 8-4 for more information about NAT and PAT support.

Dynamic PAT

This section describes dynamic PAT and includes the following topics:

- Information About Dynamic PAT, page 5-10
- Per-Session PAT vs. Multi-Session PAT (Version 9.0(1) and Later), page 5-11
- Dynamic PAT Disadvantages and Advantages, page 5-11

Information 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. If available, the real source port number is used for the mapped port. However, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 0 to 511, 512 to 1023, and 1024 to 65535. Therefore, ports below 1024 have only a small PAT pool that can be used. If you have a lot of traffic that uses the lower port ranges, you can specify a flat range of ports to be used instead of the three unequal-sized tiers.

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.

Figure 5-10 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 5-10   Dynamic PAT
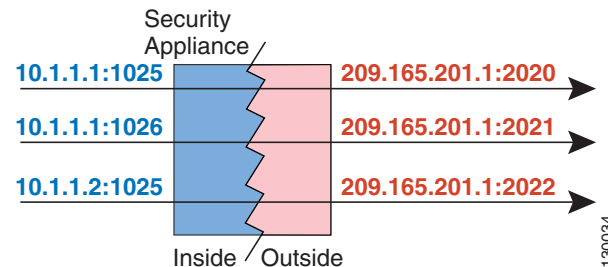
After the connection expires, the port translation also expires. For multi-session PAT, the PAT timeout is used, 30 seconds by default. For per-session PAT (9.0(1) and later), the xlate is immediately removed. Users on the destination network cannot reliably initiate a connection to a host that uses PAT (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 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.

Per-Session PAT vs. Multi-Session PAT (Version 9.0(1) and Later)

The per-session PAT feature improves the scalability of PAT and, for clustering, allows each member unit to own PAT connections; multi-session PAT connections have to be forwarded to and owned by the master unit. At the end of a per-session PAT session, the ASA sends a reset and immediately removes the xlate. This reset causes the end node to immediately release the connection, avoiding the TIME_WAIT state. Multi-session PAT, on the other hand, uses the PAT timeout, by default 30 seconds. For "hit-and-run" traffic, such as HTTP or HTTPS, the per-session feature can dramatically increase the connection rate supported by one address. Without the per-session feature, the maximum connection rate for one address for an IP protocol is approximately 2000 per second. With the per-session feature, the connection rate for one address for an IP protocol is 65535/*average-lifetime*.

By default, all TCP traffic and UDP DNS traffic use a per-session PAT xlate. For traffic that can benefit from multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT be creating a per-session deny rule. See Configuring Per-Session PAT Rules, page 6-18.

Dynamic PAT Disadvantages and Advantages

Dynamic PAT lets you use a single mapped address, thus conserving routable addresses. You can even use the ASA interface IP address as the PAT address.

Dynamic PAT does not work with some multimedia applications that have a data stream that is different from the control path. See Default Settings and NAT Limitations, page 8-4 for more information about NAT and PAT support.

Dynamic PAT may 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. (8.4(2)/8.5(1) and later) You can configure a PAT pool of addresses and use a round-robin assignment of PAT addresses to mitigate this situation.

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. Identity NAT is necessary for remote access VPN, where you need to exempt the client traffic from NAT.

Figure 5-11 shows a typical identity NAT scenario.

Figure 5-11   *Identity NAT*
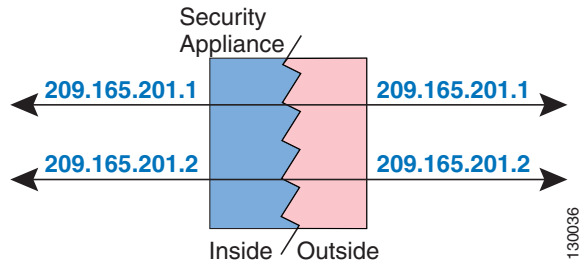
NAT in Routed and Transparent Mode

You can configure NAT in both routed and transparent firewall mode. This section describes typical usage for each firewall mode and includes the following topics:

- NAT in Routed Mode, page 5-13
- NAT in Transparent Mode, page 5-13

NAT in Routed Mode

Figure 5-12 shows a typical NAT example in routed mode, with a private network on the inside.

Figure 5-12   NAT Example: Routed Mode
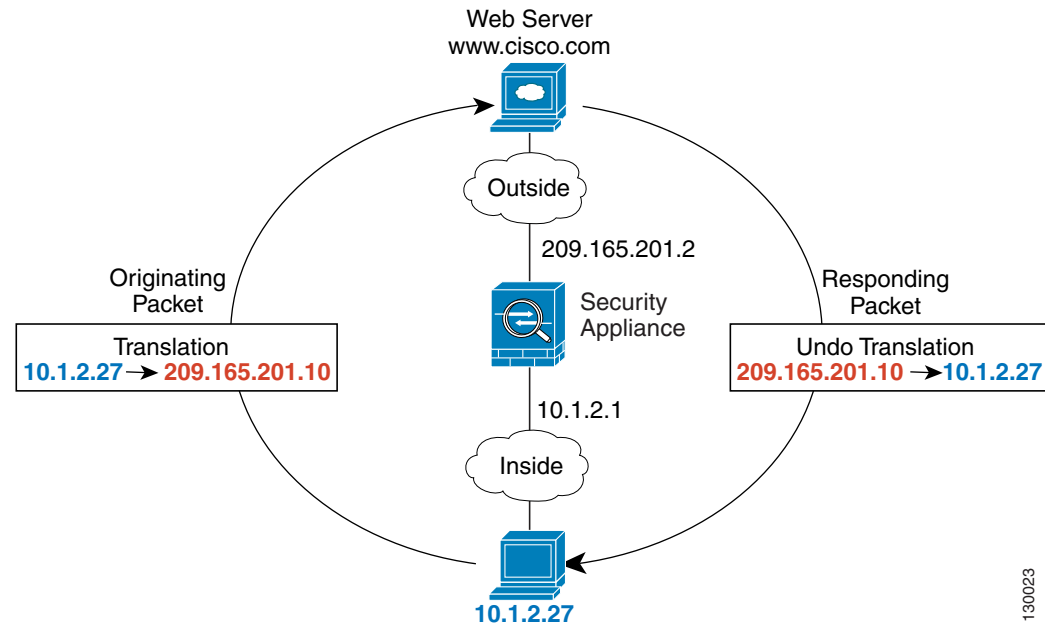
1. When the inside host at 10.1.2.27 sends a packet to a web server, the real source address of the packet, 10.1.2.27, is changed to a mapped address, 209.165.201.10.

2. When the server responds, it sends the response to the mapped address, 209.165.201.10, and the ASA receives the packet because the ASA performs proxy ARP to claim the packet.

3. The ASA 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.

NAT in Transparent Mode

Using NAT in transparent mode eliminates the need for the upstream or downstream routers to perform NAT for their networks.

NAT in transparent mode has the following requirements and limitations:

*   Because the transparent firewall does not have any interface IP addresses, you cannot use interface PAT.

*   ARP inspection is not supported. Moreover, if for some reason a host on one side of the ASA sends an ARP request to a host on the other side of the ASA, and the initiating host real address is mapped to a different address on the same subnet, then the real address remains visible in the ARP request.

*   Translating between IPv4 and IPv6 networks is not supported. Translating between two IPv6 networks, or between two IPv4 networks is supported.

Figure 5-13 shows a typical NAT scenario in transparent mode, with the same network on the inside and outside interfaces. The transparent firewall in this scenario is performing the NAT service so that the upstream router does not have to perform NAT.

Figure 5-13   *NAT Example: Transparent Mode*
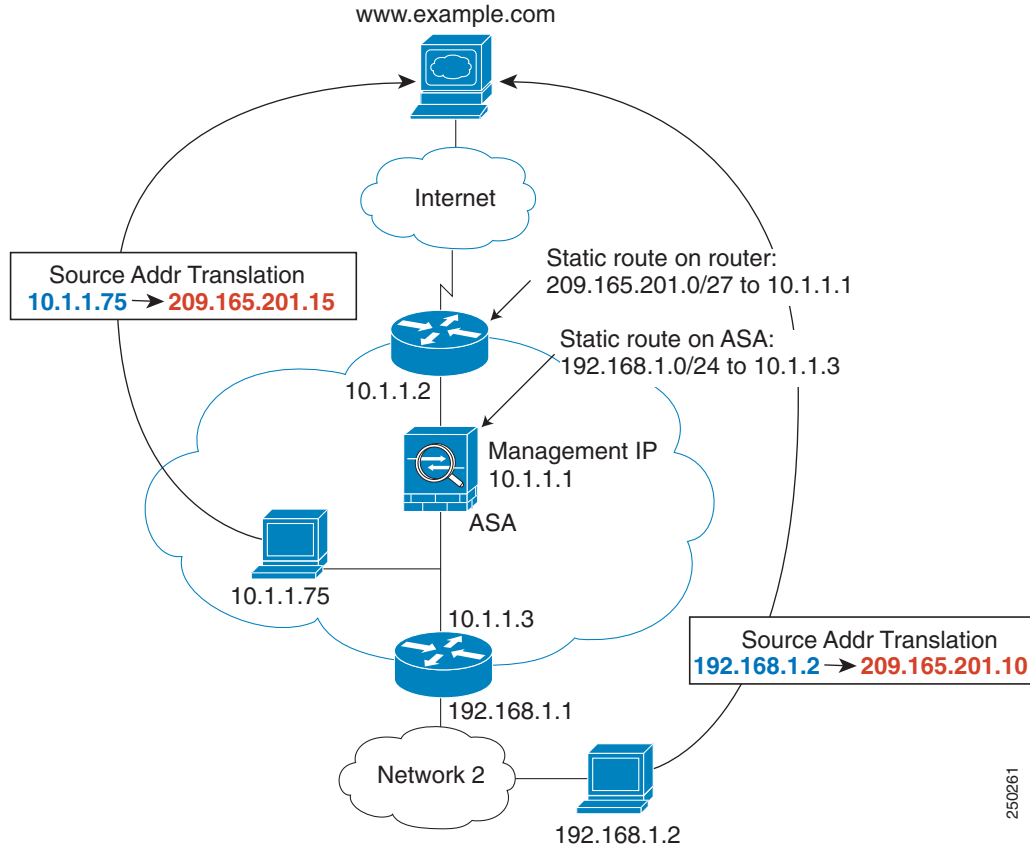
1. When the inside host at 10.1.1.75 sends a packet to a web server, the real source address of the packet, 10.1.1.75, is changed to a mapped address, 209.165.201.15.

2. When the server responds, it sends the response to the mapped address, 209.165.201.15, and the ASA receives the packet because the upstream router includes this mapped network in a static route directed to the ASA management IP address. See Mapped Addresses and Routing, page 5-22 for more information about required routes.

3. The ASA then undoes the translation of the mapped address, 209.165.201.15, back to the real address, 10.1.1.1.75. Because the real address is directly-connected, the ASA sends it directly to the host.

4. For host 192.168.1.2, the same process occurs, except for returning traffic, the ASA looks up the route in its routing table and sends the packet to the downstream router at 10.1.1.3 based on the ASA static route for 192.168.1.0/24. See Transparent Mode Routing Requirements for Remote Networks, page 5-24 for more information about required routes.

NAT and IPv6

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 (twice 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 (twice 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. You can also optionally translate the addresses net-tonet, where the first IPv4 address maps to the first IPv6 address, the second to the second, and so on.

- 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.

For specific implementation guidelines and limitations, see the configuration chapters.

How NAT is Implemented

The ASA can implement address translation in two ways: *network object NAT* and *twice NAT*. This section includes the following topics:

Main Differences Between Network Object NAT and Twice NAT

The main differences between these two NAT types are:

- How you define the real address.

  - Network object NAT—You define NAT as a parameter for a network object. A network object names an IP host, range, or subnet so you can then use the object in configuration instead of the actual IP addresses. The network object IP address serves as the real address. This method lets you easily add NAT to network objects that might already be used in other parts of your configuration.

  - Twice 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 twice NAT is more scalable.

- How source and destination NAT is implemented.
  - Network object 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.
  - Twice NAT—A single rule translates both the source and destination. A matching packet only matches the one rule, and further rules are not checked. Even if you do not configure the optional destination address for twice NAT, a matching packet still only matches one twice NAT rule. 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.
  - Network object NAT—Automatically ordered in the NAT table.
  - Twice NAT—Manually ordered in the NAT table (before or after network object NAT rules).

    See NAT Rule Order, page 5-20 for more information.

We recommend using network object NAT unless you need the extra features that twice NAT provides. Network object NAT is easier to configure, and might be more reliable for applications such as Voice over IP (VoIP). (For VoIP, because twice NAT is applicable only between two objects, you might see a failure in the translation of indirect addresses that do not belong to either of the objects.)

Information About Network Object NAT

All NAT rules that are configured as a parameter of a network object are considered to be *network object NAT* rules. Network object NAT is a quick and easy way to configure NAT for a network object, which can be a single IP address, a range of addresses, or a subnet.

After you configure the network object, you can then identify the mapped address for that object, either as an inline address or as another network object or network object group.

When a packet enters the ASA, both the source and destination IP addresses are checked against the network object 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 twice NAT for that kind of functionality (twice NAT lets you identify the source and destination address in a single rule).

To start configuring network object NAT, see Chapter 6, "Network Object NAT (ASA 8.3 and Later)."

Information About Twice NAT

Twice 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.

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.

Twice NAT also lets you use service objects for static NAT with port translation; network object NAT only accepts inline definition.

To start configuring twice NAT, see Chapter 7, "Twice NAT (ASA 8.3 and Later)."

Figure 5-14 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. When the host accesses the server at 209.165.200.225, the real address is translated to 209.165.202.130. (See Single Address for FTP, HTTP, and SMTP (Static NAT-with-Port-Translation), page 6-32 for details on how to configure this example.)

Figure 5-14   Twice NAT with Different Destination Addresses
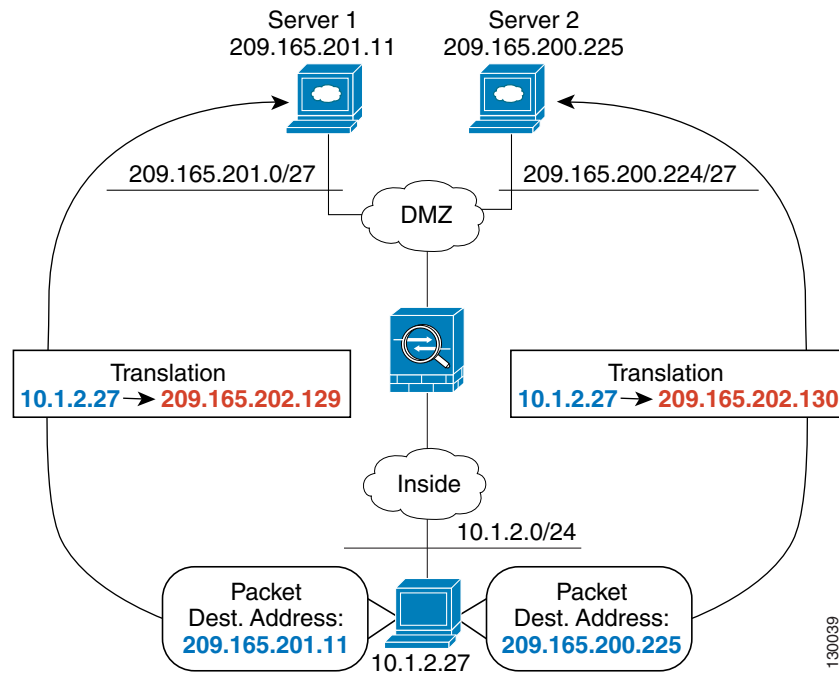
Figure 5-15 shows the use of source and destination ports. The host on the 10.1.2.0/24 network accesses a single host for both web services and Telnet services. When the host accesses the server for web services, the real address is translated to 209.165.202.129. When the host accesses the same server for Telnet services, the real address is translated to 209.165.202.130.

Figure 5-15   Twice NAT with Different Destination Ports

Figure 5-16 shows a remote host connecting to a mapped host. The mapped host has a twice static NAT translation that translates the real address only for traffic to and from the 209.165.201.0/27 network. A translation does not exist for the 209.165.200.224/27 network, so the translated host cannot connect to that network, nor can a host on that network connect to the translated host.

Figure 5-16   ***Twice Static NAT with Destination Address Translation***
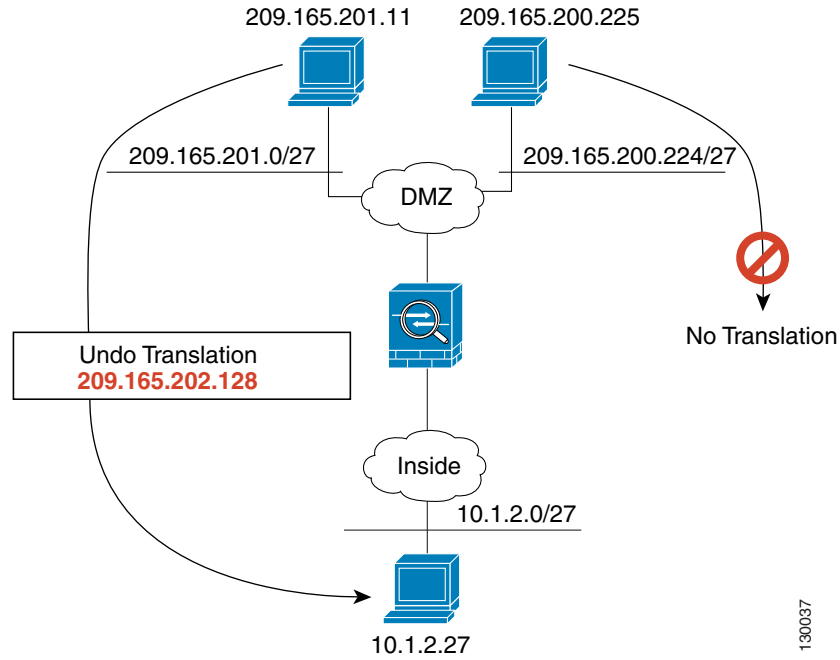
NAT Rule Order

Network object NAT rules and twice 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. Table 5-1 shows the order of rules within each section.

Table 5-1   NAT Rule Table

Table Section	Rule Type	Order of Rules within the Section
Section 1	Twice NAT	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, twice NAT rules are added to section 1.
		Note If you configure EasyVPN remote, the ASA dynamically adds invisible NAT rules to the end of this section. Be sure that you do not configure a twice NAT rule in this section that might match your VPN traffic, instead of matching the invisible rule. If VPN does not work due to NAT failure, consider adding twice NAT rules to section 3 instead.
Section 2	Network object NAT	If a match in section 1 is not found, section 2 rules are applied in the following order, as automatically determined by the ASA:
		1. Static rules.
		2. Dynamic rules.
		Within each rule type, the following ordering guidelines are used:
		a. 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.
		b. 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.
		c. 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	Twice NAT	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. You can specify whether to add a twice NAT rule to section 3 when you add the rule.

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

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 5-17).

Figure 5-17   Specifying Any Interface
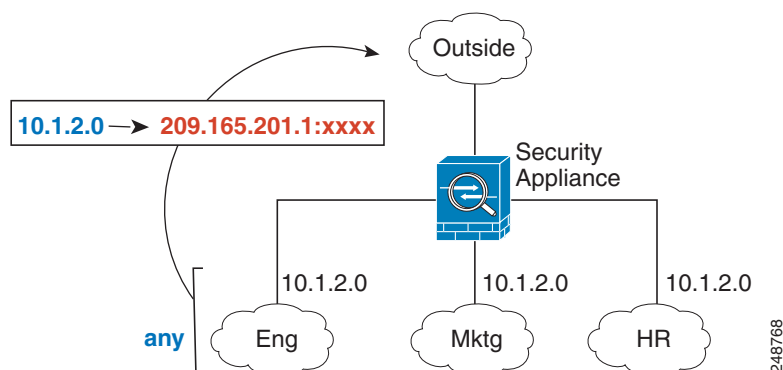
Note For transparent mode, you must choose specific source and destination interfaces.

Routing NAT Packets

The ASA needs to be the destination for any packets sent to the mapped address. The ASA also needs to determine the egress interface for any packets it receives destined for mapped addresses. This section describes how the ASA handles accepting and delivering packets with NAT, and includes the following topics:

- Mapped Addresses and Routing, page 5-22
- Transparent Mode Routing Requirements for Remote Networks, page 5-24
- Determining the Egress Interface, page 5-24

Mapped Addresses and Routing

When you translate the real address to a mapped address, the mapped address you choose determines how to configure routing, if necessary, for the mapped address.

See additional guidelines about mapped IP addresses in Chapter 6, "Network Object NAT (ASA 8.3 and Later)," and Chapter 7, "Twice NAT (ASA 8.3 and Later)."

See the following mapped address types:

- Addresses on the same network as the mapped interface.

    If you use addresses on the same network as the mapped interface, the ASA 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 ASA 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.

    Note If you configure the mapped interface to be any interface, and you specify a mapped address on the same network as one of the mapped interfaces, then if an ARP request for that mapped address comes in on a *different* interface, then you need to manually configure an ARP entry for that network on the ingress interface, specifying its MAC address (see Configuration > Device Management > Advanced > ARP > ARP Static Table). Typically, if you specify any interface for the mapped interface, then you use a unique network for the mapped addresses, so this situation would not occur.

- Addresses on a unique network.

    If you need more addresses than are available on the 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 ASA. Alternatively for routed mode, you can configure a static route on the ASA for the mapped addresses, and then redistribute the route using your routing protocol. For transparent mode, if the real host is directly-connected, configure the static route on the upstream router to point to the ASA: in 8.3, specify the global management IP address; in 8.4(1) and later, specify the bridge group IP address. For remote hosts in transparent mode, in the static route on the upstream router, you can alternatively specify the downstream router IP address.

- The same address as the real address (identity NAT).

(8.3(1), 8.3(2), and 8.4(1)) The default behavior for identity NAT has proxy ARP disabled. You cannot configure this setting.

(8.4(2) and later) The default behavior for identity NAT has proxy ARP enabled, matching other static NAT rules. You can disable proxy ARP if desired. **Note**: 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 ASA will then proxy ARP for the address, even though the packet is not actually destined for the ASA. (Note that this problem occurs even if you have a twice 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 ASA ARP response is received before the actual host ARP response, then traffic will be mistakenly sent to the ASA (see Figure 5-18).

Figure 5-18   Proxy ARP Problems with Identity NAT

In rare cases, you need proxy ARP for identity NAT; for example for virtual Telnet. When using AAA for network access, a host needs to authenticate with the ASA using a service like Telnet before any other traffic can pass. You can configure a virtual Telnet server on the ASA to provide the necessary login. When accessing the virtual Telnet address from the outside, you must configure an identity NAT rule for the address specifically for the proxy ARP functionality. Due to internal processes for virtual Telnet, proxy ARP lets the ASA keep traffic destined for the virtual Telnet address rather than send the traffic out the source interface according to the NAT rule. (See Figure 5-19).

Figure 5-19   Proxy ARP and Virtual Telnet

Transparent Mode Routing Requirements for Remote Networks

When you use NAT in transparent mode,some types of traffic require static routes. See the general operations configuration guide for more information.

Determining the Egress Interface

When the ASA receives traffic for a mapped address, the ASA unstranslates the destination address according to the NAT rule, and then it sends the packet on to the real address. The ASA determines the egress interface for the packet in the following ways:

- Transparent mode—The ASA determines the egress interface for the real address by using the NAT rule; you must specify the source and destination interfaces as part of the NAT rule.
- Routed mode—The ASA determines the egress interface in one of the following ways:
  - You configure the interface in the NAT rule—The ASA uses the NAT rule to determine the egress interface. (8.3(1) through 8.4(1)) The only exception is for identity NAT, which always uses a route lookup, regardless of the NAT configuration. (8.4(2) and later) For identity NAT, the default behavior is to use the NAT configuration. However, you have the option to always use a route lookup instead. In certain scenarios, a route lookup override is required; for example, see .
  - You do not configure the interface in the NAT rule—The ASA uses a route lookup to determine the egress interface.

Figure 5-20 shows the egress interface selection method in routed mode. In almost all cases, a route lookup is equivalent to the NAT rule interface, but in some configurations, the two methods might differ.

Figure 5-20   Routed Mode Egress Interface Selection
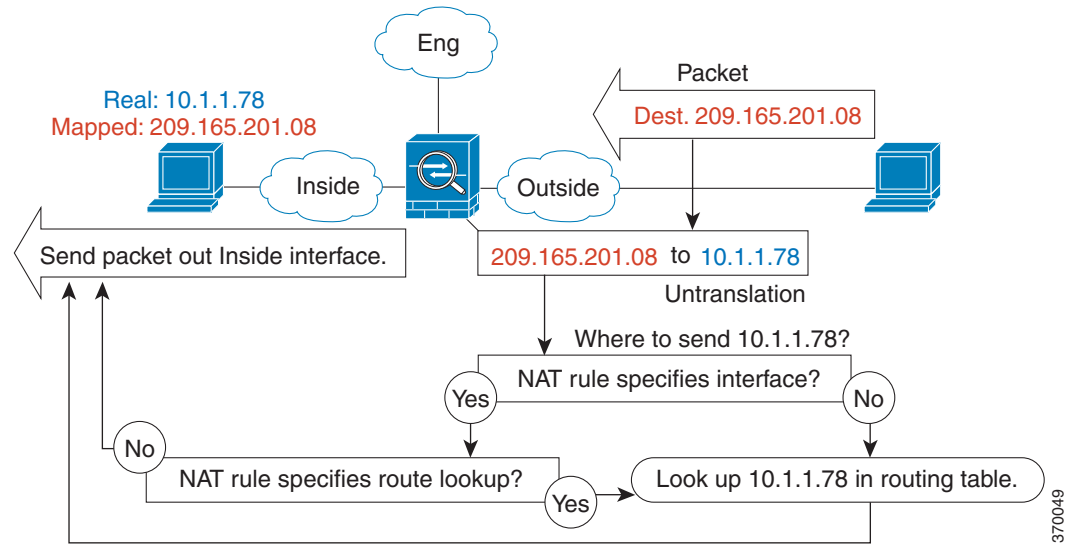
NAT for VPN

- NAT and Remote Access VPN, page 5-25
- NAT and Site-to-Site VPN, page 5-27
- NAT and VPN Management Access, page 5-29
- Troubleshooting NAT and VPN, page 5-31

NAT and Remote Access VPN

Figure 5-21 shows both an inside server (10.1.1.6) and a VPN client (209.165.201.10) accessing the Internet. Unless you configure split tunnelling for the VPN client (where only specified traffic goes through the VPN tunnel), then Internet-bound VPN traffic must also go through the ASA. When the VPN traffic enters the ASA, the ASA decrypts the packet; the resulting packet includes the VPN client local address (10.3.3.10) as the source. For both inside and VPN client local networks, you need a public IP address provided by NAT to access the Internet. The below example uses interface PAT rules. To allow the VPN traffic to exit the same interface it entered, you also need to enable intra-interface communication (AKA "hairpin" networking).

Figure 5-21   *Interface PAT for Internet-Bound VPN Traffic (Intra-Interface)*

Figure 5-22 shows a VPN client that wants to access an inside mail server. Because the ASA expects traffic between the inside network and any outside network to match the interface PAT rule you set up for Internet access, traffic from the VPN client (10.3.3.10) to the SMTP server (10.1.1.6) will be dropped due to a reverse path failure: traffic from 10.3.3.10 to 10.1.1.6 does not match a NAT rule, but returning traffic from 10.1.1.6 to 10.3.3.10 *should* match the interface PAT rule for outgoing traffic. Because forward and reverse flows do not match, the ASA drops the packet when it is received. To avoid this failure, you need to exempt the inside-to-VPN client traffic from the interface PAT rule by using an identity NAT rule between those networks. Identity NAT simply translates an address to the same address.

Figure 5-22   ***Identity NAT for VPN Clients***

See the following sample NAT configuration for the above network:

```
! Enable hairpin for non-split-tunneled VPN client traffic:
same-security-traffic permit intra-interface

! Identify local VPN network, & perform object interface PAT when going to Internet:
object network vpn_local
    subnet 10.3.3.0 255.255.255.0
    nat (outside,outside) dynamic interface

! Identify inside network, & perform object interface PAT when going to Internet:
object network inside_nw
    subnet 10.1.1.0 255.255.255.0
    nat (inside,outside) dynamic interface

! Use twice NAT to pass traffic between the inside network and the VPN client without
! address translation (identity NAT):
nat (inside,outside) source static inside_nw inside_nw destination static vpn_local
vpn_local
```

NAT and Site-to-Site VPN

Figure 5-23 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 5-23   Interface PAT and Identity NAT for Site-to-Site VPN
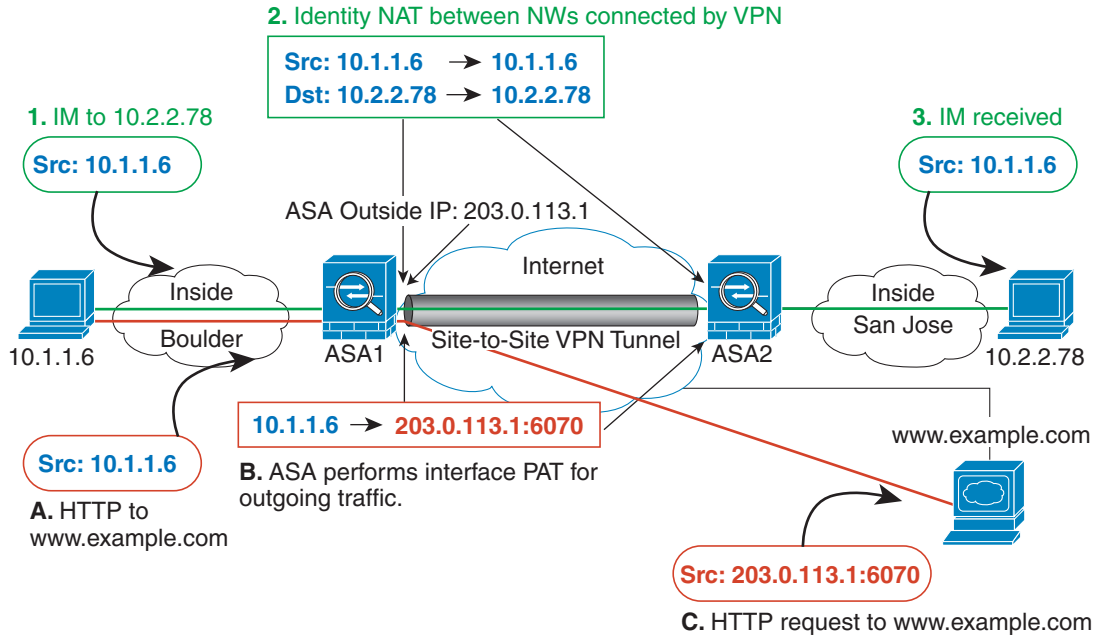
Figure 5-24 shows a VPN client connected to ASA1 (Boulder), with a Telnet request for a server (10.2.2.78) accessible over a site-to-site tunnel between ASA1 and ASA2 (San Jose). Because this is a hairpin connection, you need to enable intra-interface communication, which is also required for non-split-tunneled Internet-bound traffic from the VPN client. You also need to configure identity NAT between the VPN client and the Boulder & San Jose networks, just as you would between any networks connected by VPN to exempt this traffic from outbound NAT rules.

Figure 5-24   VPN Client Access to Site-to-Site VPN
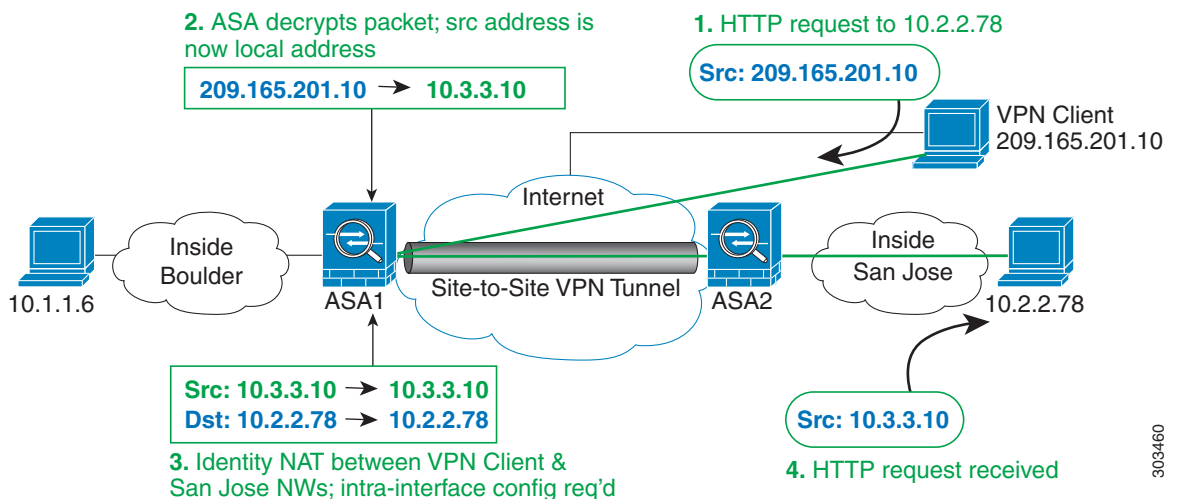
See the following sample NAT configuration for ASA1 (Boulder):

```
! Enable hairpin for VPN client traffic:
same-security-traffic permit intra-interface

! Identify local VPN network, & perform object interface PAT when going to Internet:
```

```
object network vpn_local
    subnet 10.3.3.0 255.255.255.0
    nat (outside,outside) dynamic interface

! Identify inside Boulder network, & perform object interface PAT when going to Internet:
object network boulder_inside
    subnet 10.1.1.0 255.255.255.0
    nat (inside,outside) dynamic interface

! Identify inside San Jose network for use in twice NAT rule:
object network sanjose_inside
    subnet 10.2.2.0 255.255.255.0

! Use twice NAT to pass traffic between the Boulder network and the VPN client without
! address translation (identity NAT):
nat (inside,outside) source static boulder_inside boulder_inside destination static
vpn_local vpn_local

! Use twice NAT to pass traffic between the Boulder network and San Jose without
! address translation (identity NAT):
nat (inside,outside) source static boulder_inside boulder_inside destination static
sanjose_inside sanjose_inside

! Use twice NAT to pass traffic between the VPN client and San Jose without
! address translation (identity NAT):
nat (outside,outside) source static vpn_local vpn_local destination static sanjose_inside
sanjose_inside
```

See the following sample NAT configuration for ASA2 (San Jose):

```
! Identify inside San Jose network, & perform object interface PAT when going to Internet:
object network sanjose_inside
    subnet 10.2.2.0 255.255.255.0
    nat (inside,outside) dynamic interface

! Identify inside Boulder network for use in twice NAT rule:
object network boulder_inside
    subnet 10.1.1.0 255.255.255.0

! Identify local VPN network for use in twice NAT rule:
object network vpn_local
    subnet 10.3.3.0 255.255.255.0

! Use twice NAT to pass traffic between the San Jose network and Boulder without
! address translation (identity NAT):
nat (inside,outside) source static sanjose_inside sanjose_inside destination static
boulder_inside boulder_inside

! Use twice NAT to pass traffic between the San Jose network and the VPN client without
! address translation (identity NAT):
nat (inside,outside) source static sanjose_inside sanjose_inside destination static
vpn_local vpn_local
```

NAT and VPN Management Access

When using VPN, you can allow management access to an interface other than the one from which you entered the ASA. For example, if you enter the ASA from the outside interface, the management-access feature lets you connect to the inside interface using ASDM, SSH, Telnet, or SNMP; or you can ping the inside interface.

Figure 5-25 shows a VPN client Telnetting to the ASA inside interface. When you use a management-access interface, and you configure identity NAT according to the NAT and Remote Access VPN, page 5-25 or NAT and Site-to-Site VPN, page 5-27 section, you must configure NAT with the route lookup option. Without route lookup, the ASA sends traffic out the interface specified in the NAT command, regardless of what the routing table says; in the below example, the egress interface is the inside interface. You do not want the ASA to send the management traffic out to the inside network; it will never return to the inside interface IP address. The route lookup option lets the ASA send the traffic directly to the inside interface IP address instead of to the inside network. For traffic from the VPN client to a host on the inside network, the route lookup option will still result in the correct egress interface (inside), so normal traffic flow is not affected. See the Determining the Egress Interface, page 5-24 for more information about the route lookup option.

Figure 5-25   **VPN Management Access**
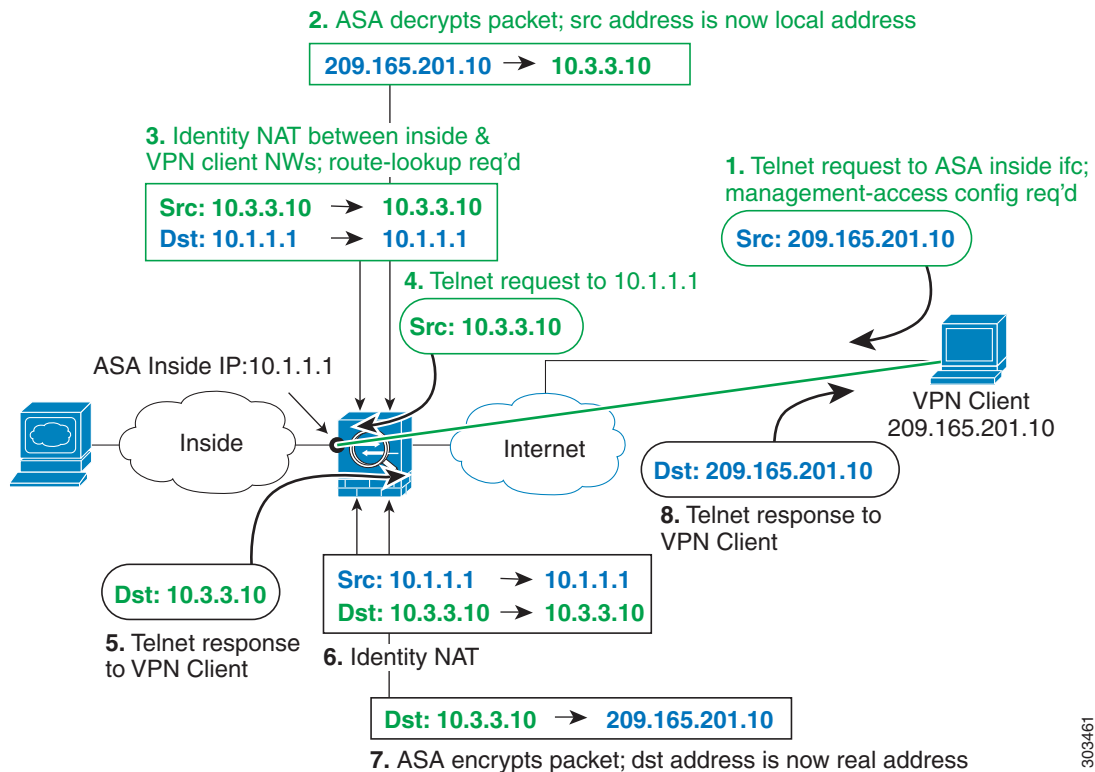
See the following sample NAT configuration for the above network:

```
! Enable hairpin for non-split-tunneled VPN client traffic:
same-security-traffic permit intra-interface

! Enable management access on inside ifc:
management-access inside

! Identify local VPN network, & perform object interface PAT when going to Internet:
object network vpn_local
    subnet 10.3.3.0 255.255.255.0
    nat (outside,outside) dynamic interface

! Identify inside network, & perform object interface PAT when going to Internet:
object network inside_nw
    subnet 10.1.1.0 255.255.255.0
    nat (inside,outside) dynamic interface
```

```
! Use twice NAT to pass traffic between the inside network and the VPN client without
! address translation (identity NAT), w/route-lookup:
nat (outside,inside) source static vpn_local vpn_local destination static inside_nw
inside_nw route-lookup
```

Troubleshooting NAT and VPN

See the following monitoring tools for troubleshooting NAT issues with VPN:

- Packet tracer—When used correctly, a packet tracer shows which NAT rules a packet is hitting.
- **show nat detail**—Shows hit counts and untranslated traffic for a given NAT rule.
- **show conn all**—Lets you see active connections including to and from the box traffic.

To familiarize yourself with a non-working configuration vs. a working configuration, you can perform the following steps:

1. Configure VPN without identity NAT.
2. Enter **show nat detail** and **show conn all**.
3. Add the identity NAT configuration.
- Repeat **show nat detail** and **show conn all**.

DNS and NAT

You might need to configure the ASA 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.

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.

Note   DNS rewrite is not applicable for PAT because multiple PAT rules are applicable for each A-record, and the PAT rule to use is ambiguous.

Note   If you configure a twice NAT rule, you cannot configure DNS modification if you specify the source address as well as the destination address. These kinds of rules can potentially have a different translation for a single address when going to A vs. B. Therefore, the ASA cannot accurately match the IP address inside the DNS reply to the correct twice NAT rule; the DNS reply does not contain information about which source/destination address combination was in the packet that prompted the DNS request.

Note   This feature requires DNS application inspection to be enabled, which it is by default. See DNS Inspection, page 9-1 for more information.

Figure 5-26 shows a DNS server that is accessible from the outside interface. A server, ftp.cisco.com, is on the inside interface. You configure the ASA 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 ASA 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.

Figure 5-26   DNS Reply Modification, DNS Server on Outside
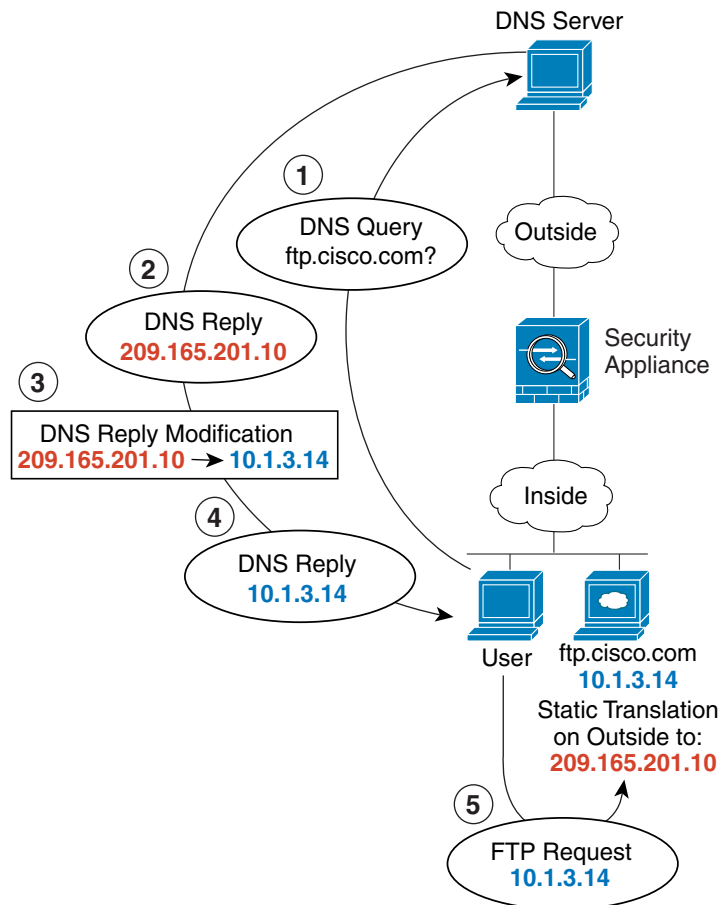
Figure 5-27 shows a user on the inside network requesting the IP address for ftp.cisco.com, which is on the DMZ network, from an outside DNS server. The DNS server replies with the mapped address (209.165.201.10) according to the static rule between outside and DMZ even though the user is not on the DMZ network. The ASA translates the address inside the DNS reply to 10.1.3.14. If the user needs to access ftp.cisco.com using the real address, then no further configuration is required. If there is also

a static rule between the inside and DMZ, then you also need to enable DNS reply modification on this rule. The DNS reply will then be modified two times.In this case, the ASA again translates the address inside the DNS reply to 192.168.1.10 according to the static rule between inside and DMZ.

Figure 5-27   DNS Reply Modification, DNS Server, Host, and Server on Separate Networks
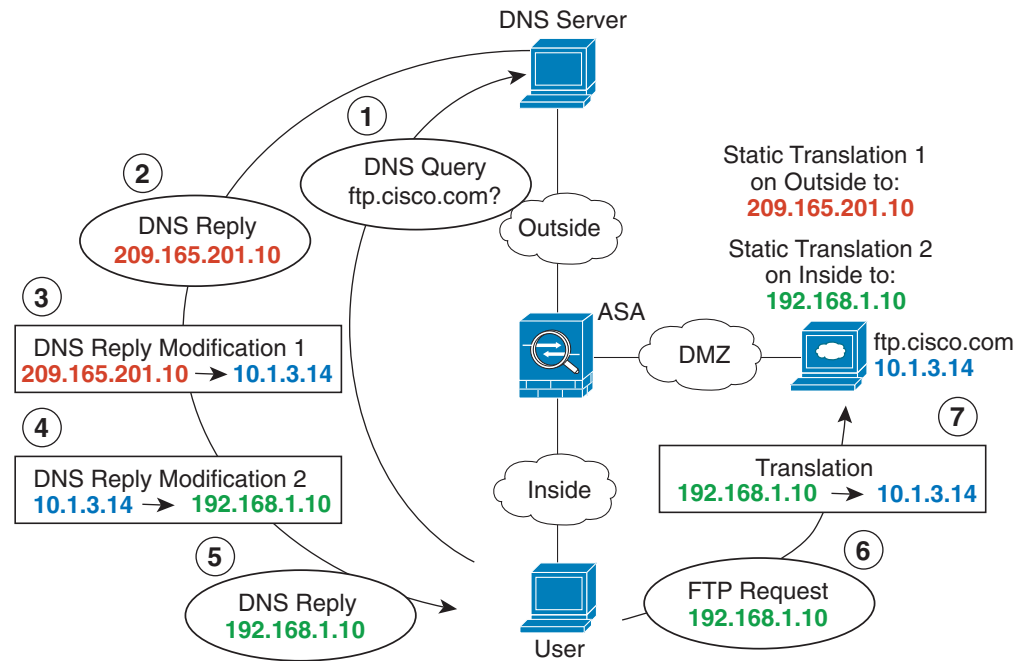
Figure 5-28 shows an FTP server and DNS server on the outside. The ASA has a static translation for the outside server. In this case, when an inside user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.20.10. Because you want inside users to use the mapped address for ftp.cisco.com (10.1.2.56) you need to configure DNS reply modification for the static translation.

Figure 5-28   DNS Reply Modification, DNS Server on Host Network
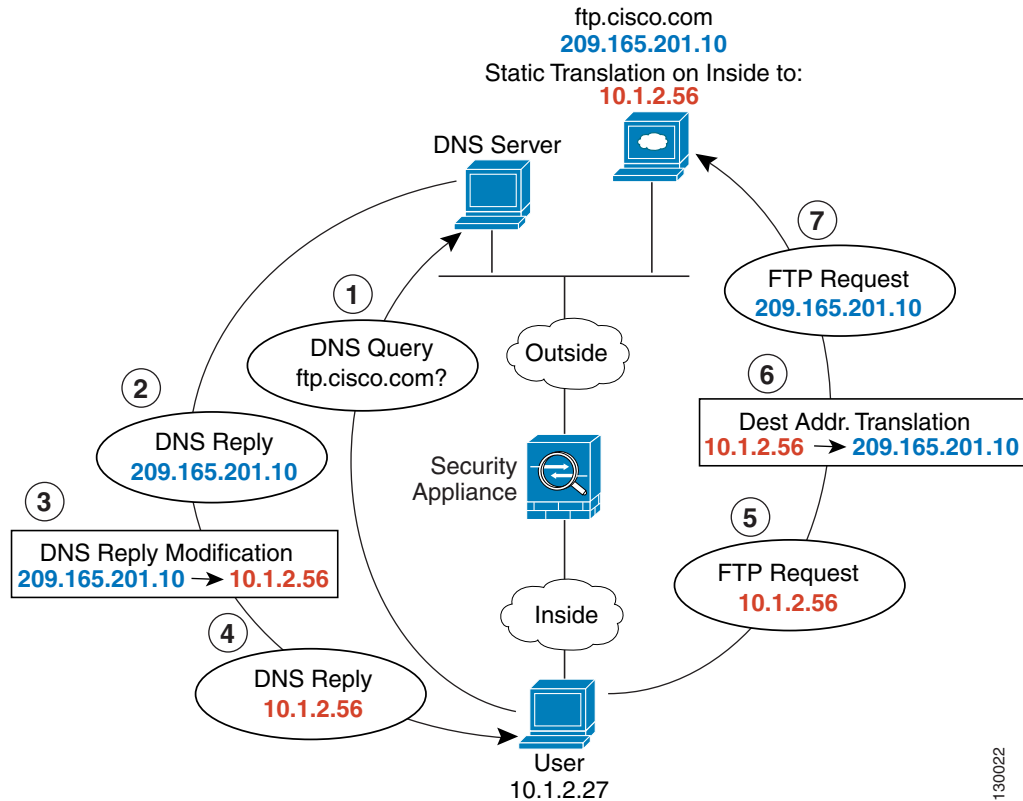
Figure 5-28 shows an FTP server and DNS server on the outside IPv4 network. The ASA has a static translation for the outside server. In this case, when an inside IPv6 user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.200.225.

Because you want inside users to use the mapped address for ftp.cisco.com (2001:DB8::D1A5:C8E1) 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.

Figure 5-29   *DNS64 Reply Modification Using Outside NAT*
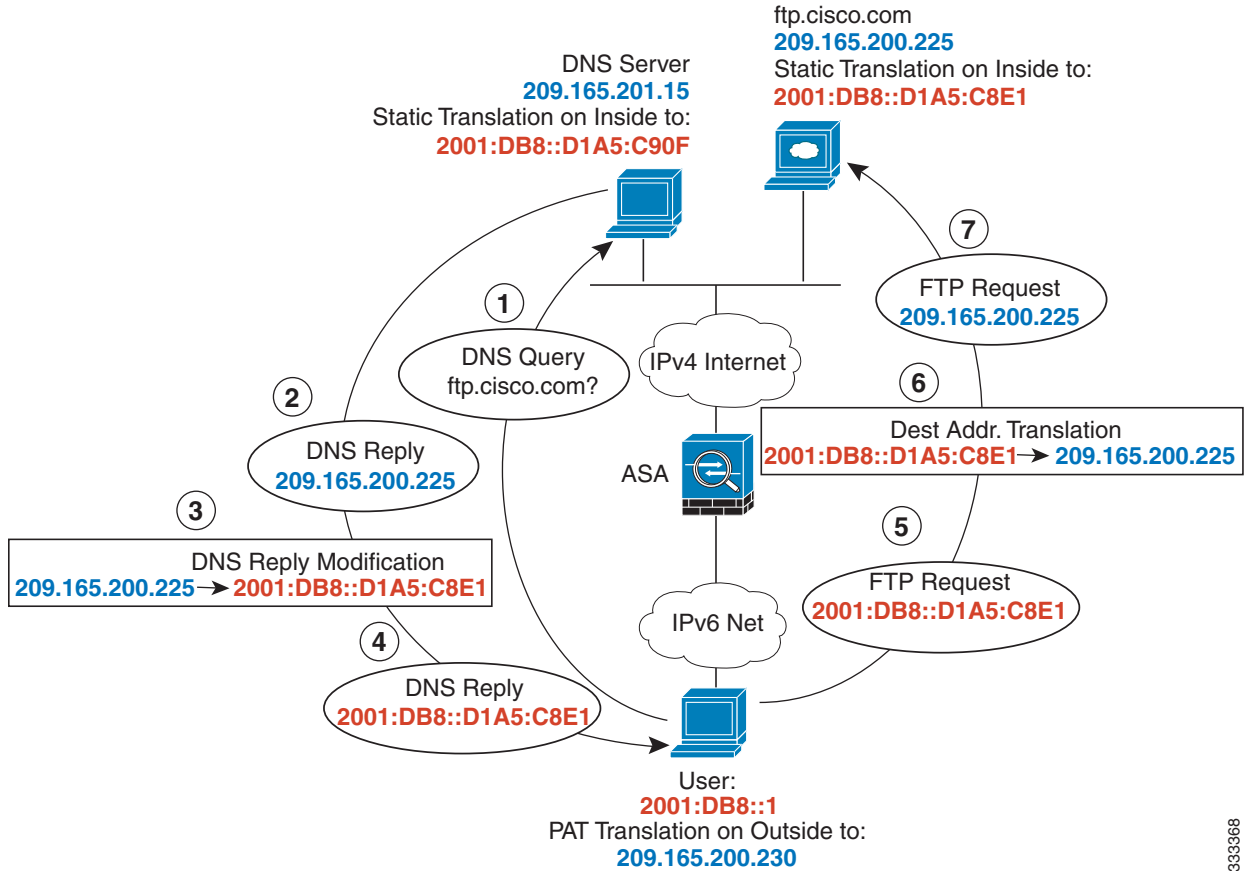
Figure 5-30 shows an FTP server and DNS server on the outside. The ASA has a static translation for the outside server. In this case, when an inside user performs a reverse DNS lookup for 10.1.2.56, the ASA modifies the reverse DNS query with the real address, and the DNS server responds with the server name, ftp.cisco.com.

Figure 5-30   PTR Modification, DNS Server on Host Network
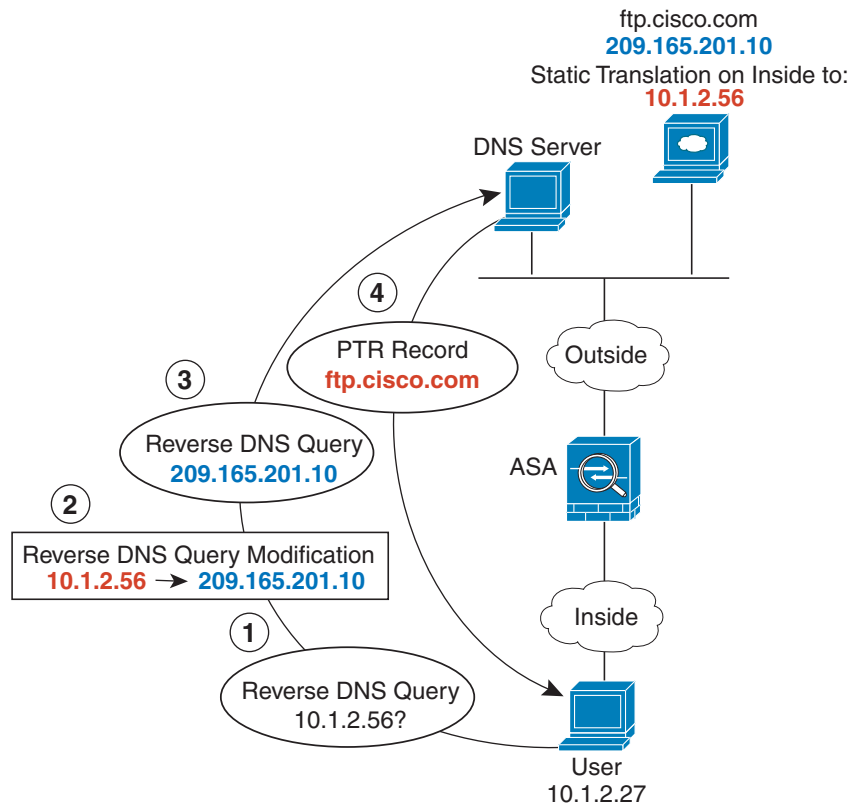
Where to Go Next

To configure network object NAT, see Chapter 6, "Network Object NAT (ASA 8.3 and Later)."

To configure twice NAT, see Chapter 7, "Twice NAT (ASA 8.3 and Later)."

CHAPTER 6

Network Object NAT (ASA 8.3 and Later)

All NAT rules that are configured as a parameter of a network object are considered to be *network object NAT* rules. Network object NAT is a quick and easy way to configure NAT for a single IP address, a range of addresses, or a subnet. After you configure the network object, you can then identify the mapped address for that object.

This chapter describes how to configure network object NAT, and it includes the following sections:

- Information About Network Object NAT, page 6-1
- Licensing Requirements for Network Object NAT, page 6-2
- Prerequisites for Network Object NAT, page 6-2
- Guidelines and Limitations, page 6-2
- Default Settings, page 6-3
- Configuring Network Object NAT, page 6-4
- Monitoring Network Object NAT, page 6-19
- Configuration Examples for Network Object NAT, page 6-20
- Feature History for Network Object NAT, page 6-45

Note For detailed information about how NAT works, see Chapter 5, "Information About NAT (ASA 8.3 and Later)."

Information About Network Object NAT

When a packet enters the ASA, both the source and destination IP addresses are checked against the network object 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 a source address should be translated to A when going to destination X, but be translated to B when going to destination Y. Use twice NAT for that kind of functionality (twice NAT lets you identify the source and destination address in a single rule).

For detailed information about the differences between twice NAT and network object NAT, see How NAT is Implemented, page 5-15.

Network object NAT rules are added to section 2 of the NAT rules table. For more information about NAT ordering, see NAT Rule Order, page 5-20.

Licensing Requirements for Network Object NAT

The following table shows the licensing requirements for this feature:

Model	License Requirement
ASAv	Standard or Premium License.
All other models	Base License.

Prerequisites for Network Object NAT

Depending on the configuration, you can configure the mapped address inline if desired or you can create a separate network object or network object group for the mapped address. Network object groups are particularly useful for creating a mapped address pool with discontinous IP address ranges or multiple hosts or subnets. To create a network object or group, see the general operations configuration guide.

For specific guidelines for objects and groups, see the configuration section for the NAT type you want to configure. See also the Guidelines and Limitations, page 6-2 section.

Guidelines and Limitations

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

- Supported in routed and transparent firewall mode.

- In transparent mode, you must specify the real and mapped interfaces; you cannot use --Any--.

- In transparent mode, you cannot configure interface PAT, because the transparent mode interfaces do not have IP addresses. You also cannot use the management IP address as a mapped address.

- In transparent mode, translating between IPv4 and IPv6 networks is not supported. Translating between two IPv6 networks, or between two IPv4 networks is supported.

IPv6 Guidelines

- Supports IPv6. See also the NAT and IPv6, page 5-15.

- For routed mode, you can also translate between IPv4 and IPv6.

- For transparent mode, translating between IPv4 and IPv6 networks is not supported. Translating between two IPv6 networks, or between two IPv4 networks is supported.

- For transparent mode, a PAT pool is not supported for IPv6.

- 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.

Additional Guidelines

- 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, for example, **object network obj-10.10.10.1-01**, **object network obj-10.10.10.1-02**, and so on.

- 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. However, clearing the translation table disconnects all current connections that use translations.

    > **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** command. This safeguard ensures that the same address is not assigned to multiple hosts.

- Objects and object groups used in NAT cannot be undefined; they must include IP addresses.

- You cannot use an object group with both IPv4 and IPv6 addresses; the object group must include only one type of address.

- 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), use the interface name instead of the IP address.
    - (Transparent mode) The management IP address.
    - (Dynamic NAT) The standby interface IP address when VPN is enabled.
    - Existing VPN pool addresses.

- For application inspection limitations with NAT or PAT, see Default Settings and NAT Limitations, page 8-4 in Chapter 8, "Getting Started with Application Layer Protocol Inspection."

Default Settings

- (Routed mode) The default real and mapped interface is Any, which applies the rule to all interfaces.

- (8.3(1), 8.3(2), and 8.4(1)) The default behavior for identity NAT has proxy ARP disabled. You cannot configure this setting. (8.4(2) and later) The default behavior for identity NAT has proxy ARP enabled, matching other static NAT rules. You can disable proxy ARP if desired. See Routing NAT Packets, page 5-22 for more information.

- If you specify an optional interface, then the ASA uses the NAT configuration to determine the egress interface. (8.3(1) through 8.4(1)) The only exception is for identity NAT, which always uses a route lookup, regardless of the NAT configuration. (8.4(2) and later) For identity NAT, the default behavior is to use the NAT configuration, but you have the option to always use a route lookup instead. See Routing NAT Packets, page 5-22 for more information.

Configuring Network Object NAT

This section describes how to configure network object NAT and includes the following topics:

Configuring Dynamic NAT or Dynamic PAT Using a PAT Pool

This section describes how to configure network object NAT for dynamic NAT or for dynamic PAT using a PAT pool. For more information, see Dynamic NAT, page 5-8 or the Dynamic PAT, page 5-10.

Guidelines

For a PAT pool:

- If available, the real source port number is used for the mapped port. However, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 0 to 511, 512 to 1023, and 1024 to 65535. Therefore, ports below 1024 have only a small PAT pool that can be used. (8.4(3) and later, not including 8.5(1) or 8.6(1)) If you have a lot of traffic that uses the lower port ranges, you can now specify for a PAT pool a flat range of ports to be used instead of the three unequal-sized tiers: either 1024 to 65535, or 1 to 65535.

- If you use the same PAT pool object in two separate rules, then be sure to specify the same options for each rule. For example, if one rule specifies extended PAT and a flat range, then the other rule must also specify extended PAT and a flat range.

For extended PAT for a PAT pool:

- Many application inspections do not support extended PAT. See Default Settings and NAT Limitations, page 8-4 in Chapter 8, "Getting Started with Application Layer Protocol Inspection," for a complete list of unsupported inspections.

- If you enable extended PAT for a dynamic PAT rule, then you cannot also use an address in the PAT pool as the PAT address in a separate static NAT with port translation rule. For example, if the PAT pool includes 10.1.1.1, then you cannot create a static NAT-with-port-translation rule using 10.1.1.1 as the PAT address.

- If you use a PAT pool and specify an interface for fallback, you cannot specify extended PAT.

- For VoIP deployments that use ICE or TURN, do not use extended PAT. ICE and TURN rely on the PAT binding to be the same for all destinations.

For round robin for a PAT pool:

- If a host has an existing connection, then subsequent connections from that host will use the same PAT IP address if ports are available. **Note**: This "stickiness" does not survive a failover. If the ASA fails over, then subsequent connections from a host may not use the initial IP address.

- Round robin, especially when combined with extended PAT, can consume a large amount of memory. Because NAT pools are created for every mapped protocol/IP address/port range, round robin results in a large number of concurrent NAT pools, which use memory. Extended PAT results in an even larger number of concurrent NAT pools.

Detailed Steps

Step 1   Add NAT to a new or existing network object:

- To add a new network object, choose **Configuration > Firewall > NAT Rules**, then click **Add > Add Network Object NAT Rule**.

- To add NAT to an existing network object, choose **Configuration > Firewall > Objects > Network Objects/Groups**, and then double-click a network object.

For more information, see the general operations configuration guide.

The Add/Edit Network Object dialog box appears.

Step 2   For a new object, enter values for the following fields:

a.   Name—The object name. Use characters a to z, A to Z, 0 to 9, a period, a dash, a comma, or an underscore. The name must be 64 characters or less.

b.   Type—Host, Network, or Range.

c.   IP Address—An IPv4 or IPv6 address. If you select Range as the object type, the IP Address field changes to allow you to enter a Start Address and an End address.

d.   Netmask/Prefix Length—Enter the subnet mask or prefix length.

e.   Description—(Optional) The description of the network object (up to 200 characters in length).

Step 3   If the NAT section is hidden, click **NAT** to expand the section.

Step 4   Check the **Add Automatic Translation Rules** check box.

Step 5   From the Type drop-down list, choose **Dynamic**. Choose **Dynamic** even if you are configuring dynamic PAT with a PAT pool.

Step 6   Configure either dynamic NAT, or dynamic PAT with a PAT pool:

- Dynamic NAT—To the right of the Translated Addr field, click the browse button and choose an existing network object or create a new object from the Browse Translated Addr dialog box.

Note   The object or group cannot contain a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

- Dynamic PAT using a PAT pool—Enable a PAT pool:

a. Do not enter a value for the Translated Addr. field; leave it blank.

b. Check the **PAT Pool Translated Address** check box, then click the browse button and choose an existing network object or create a new network object from the Browse Translated PAT Pool Address dialog box.

✎

Note   The PAT pool object or group cannot contain a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

c. (Optional) Check the **Round Robin** check box to assign addresses/ports in a round-robin fashion. By default without round robin, all ports for a PAT address will be allocated before the next PAT address is used. The round-robin method assigns one address/port from each PAT address in the pool before returning to use the first address again, and then the second address, and so on.

d. (Optional, 8.4(3) and later, not including 8.5(1) or 8.6(1)) Check the **Extend PAT uniqueness to per destination instead of per interface** check box to use extended PAT. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80.

e. (Optional, 8.4(3) and later, not including 8.5(1) or 8.6(1)) Check the **Translate TCP or UDP ports into flat range (1024-65535)** check box to use the 1024 to 65535 port range as a single flat range when allocating ports. When choosing the mapped port number for a translation, the ASA uses the real source port number if it is available. However, without this option, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 1 to 511, 512 to 1023, and 1024 to 65535. To avoid running out of ports at the low ranges, configure this setting. To use the entire range of 1 to 65535, also check the **Include range 1 to 1023** check box.

Step 7   (Optional, Routed Mode Only) To use the interface IP address as a backup method when the other mapped addresses are already allocated, check the **Fall through to interface PAT (dest intf)** check box, and choose the interface from the drop-down list. To use the IPv6 address of the interface, also check the **Use IPv6 for interface PAT** checkbox.

Step 8   (Optional) Click **Advanced**, and configure the following options in the Advanced NAT Settings dialog box.

- Translate DNS replies for rule—Translates the IP address in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). See DNS and NAT, page 5-31 for more information.

- (Required for Transparent Firewall Mode) Source Interface—Specifies the real interface where this NAT rule applies. By default, the rule applies to all interfaces.

- (Required for Transparent Firewall Mode) Destination Interface—Specifies the mapped interface where this NAT rule applies. By default, the rule applies to all interfaces.

When you are finished, click **OK**. You return to the Add/Edit Network Object dialog box.

Step 9   Click **OK**, and then **Apply**.

Configuring Dynamic PAT (Hide)

This section describes how to configure network object NAT for dynamic PAT (hide). For dynamic PAT using a PAT pool, see Configuring Dynamic NAT or Dynamic PAT Using a PAT Pool, page 6-4 instead of using this section. For more information, see Dynamic PAT, page 5-10.

Detailed Steps

Step 1   Add NAT to a new or existing network object:

- To add a new network object, choose **Configuration > Firewall > NAT Rules**, then click **Add > Add Network Object NAT Rule**.

- To add NAT to an existing network object, choose **Configuration > Firewall > Objects > Network Objects/Groups**, and then double-click a network object.

For more information, see the general operations configuration guide.

The Add/Edit Network Object dialog box appears.

Step 2 For a new object, enter values for the following fields:

  a. Name—The object name. Use characters a to z, A to Z, 0 to 9, a period, a dash, a comma, or an underscore. The name must be 64 characters or less.

  b. Type—Host, Network, or Range.

  c. IP Address—An IPv4 or IPv6 address. If you select Range as the object type, the IP Address field changes to allow you to enter a Start Address and an End address.

  d. Netmask/Prefix Length—Enter the subnet mask or prefix length.

  e. Description—(Optional) The description of the network object (up to 200 characters in length).

Step 3 If the NAT section is hidden, click **NAT** to expand the section.

Step 4   Check the **Add Automatic Translation Rules** check box.

Step 5   From the Type drop-down list, choose **Dynamic PAT (Hide)**.

> ✎
> **Note**   To configure dynamic PAT using a PAT pool instead of a single address, see Configuring
> Dynamic NAT or Dynamic PAT Using a PAT Pool, page 6-4.

Step 6   Specify a single mapped address. In the Translated Addr. field, specify the mapped IP address by doing
one of the following:

- Type a host IP address.
- Type an interface name or click the browse button, and choose an interface from the Browse
  Translated Addr dialog box.

If you specify an interface name, then you enable *interface PAT*, where the specified interface IP
address is used as the mapped address. To use the IPv6 interface address, you must also check the
Use IPv6 for interface PAT checkbox. With interface PAT, the NAT rule only applies to the
specified mapped interface. (If you do not use interface PAT, then the rule applies to all interfaces
by default.) See Step 7 to optionally also configure the real interface to be a specific interface instead
of --Any--.

Note   You cannot specify an interface in transparent mode.

- Click the browse button, and choose an existing host address from the Browse Translated Addr dialog box.
- Click the browse button, and create a new named object from the Browse Translated Addr dialog box.

Step 7   (Optional) Click **Advanced**, and configure the following options in the Advanced NAT Settings dialog box.

- Translate DNS replies for rule—Translates the IP address in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). See DNS and NAT, page 5-31 for more information.
- (Required for Transparent Firewall Mode) Source Interface—Specifies the real interface where this NAT rule applies. By default, the rule applies to all interfaces.
- (Required for Transparent Firewall Mode) Destination Interface—Specifies the mapped interface where this NAT rule applies. By default, the rule applies to all interfaces.

When you are finished, click **OK**. You return to the Add/Edit Network Object dialog box.

Step 8   Click **OK**, and then **Apply**.

Configuring Static NAT or Static NAT-with-Port-Translation

This section describes how to configure a static NAT rule using network object NAT. For more information, see Static NAT, page 5-3.

Detailed Steps

Step 1   Add NAT to a new or existing network object:

- To add a new network object, choose **Configuration > Firewall > NAT Rules**, then click **Add > Add Network Object NAT Rule**.

- To add NAT to an existing network object, choose **Configuration > Firewall > Objects > Network Objects/Groups**, and then double-click a network object.

For more information, see the general operations configuration guide.

The Add/Edit Network Object dialog box appears.

Step 2   For a new object, enter values for the following fields:

   a.   Name—The object name. Use characters a to z, A to Z, 0 to 9, a period, a dash, a comma, or an underscore. The name must be 64 characters or less.

   b.   Type—Network, Host, or Range.

   c.   IP Address—An IPv4 or IPv6 address. If you select Range as the object type, the IP Address field changes to allow you to enter a Start Address and an End address.

   d.   Netmask/Prefix Length—Enter the subnet mask or prefix length.

   e.   Description—(Optional) The description of the network object (up to 200 characters in length).

Step 3   If the NAT section is hidden, click **NAT** to expand the section.

Step 4   Check the **Add Automatic Translation Rules** check box.

Step 5   From the Type drop-down list, choose **Static**.

Step 6   In the Translated Addr. field, do one of the following:

- Type an IP address.

  When you type an IP address, the netmask or range for the mapped network is the same as that of the real network. For example, if the real network is a host, then this address will be a host address. In the case of a range, then the mapped addresses include the same number of addresses as the real range. For example, if the real address is defined as a range from 10.1.1.1 through 10.1.1.6, and you specify 172.20.1.1 as the mapped address, then the mapped range will include 172.20.1.1 through 172.20.1.6.

- (For static NAT-with-port-translation only) Type an interface name or click the browse button, and choose an interface from the Browse Translated Addr dialog box.

  To use the IPv6 interface address, you must also check the **Use IPv6 for interface PAT** checkbox. Be sure to also configure a service on the Advanced NAT Settings dialog box (see Step 8). (You cannot specify an interface in transparent mode).

- Click the browse button, and choose an existing address from the Browse Translated Addr dialog box.

- Click the browse button, and create a new address from the Browse Translated Addr dialog box.

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. For more information, see Static NAT, page 5-3.

Step 7   (Optional) For NAT46, check **Use one-to-one address translation**. For NAT 46, specify one-to-one to translate the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used. For a one-to-one translation, you must use this keyword.

Step 8   (Optional) Click **Advanced**, and configure the following options in the Advanced NAT Settings dialog box.

- Translate DNS replies for rule—Translates the IP address in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). See DNS and NAT, page 5-31 for more information.

- Disable Proxy ARP on egress interface—Disables proxy ARP for incoming packets to the mapped IP addresses. See Mapped Addresses and Routing, page 5-22 for more information.

- (Required for Transparent Firewall Mode) Interface:
    - Source Interface—Specifies the real interface where this NAT rule applies. By default, the rule applies to all interfaces.
    - Destination Interface—Specifies the mapped interface where this NAT rule applies. By default, the rule applies to all interfaces.

- Service:
    - Protocol—Configures static NAT-with-port-translation. Choose **tcp** or **udp**.
    - Real Port—You can type either a port number or a well-known port name (such as "ftp").
    - Mapped Port—You can type either a port number or a well-known port name (such as "ftp").

When you are finished, click **OK**. You return to the Add/Edit Network Object dialog box.

Step 9   Click **OK**, and then **Apply**.

Because static rules are bidirectional (allowing initiation to and from the real host), the NAT Rules table show two rows for each static rule, one for each direction.

Configuring Identity NAT

This section describes how to configure an identity NAT rule using network object NAT. For more information, see .

Detailed Steps

Step 1   Add NAT to a new or existing network object:

- To add a new network object, choose **Configuration > Firewall > NAT Rules**, then click **Add > Add Network Object NAT Rule**.

- To add NAT to an existing network object, choose **Configuration > Firewall > Objects > Network Objects/Groups**, and then double-click a network object.

For more information, see the general operations configuration guide.

The Add/Edit Network Object dialog box appears.

Step 2   For a new object, enter values for the following fields:

a.   Name—The object name. Use characters a to z, A to Z, 0 to 9, a period, a dash, a comma, or an underscore. The name must be 64 characters or less.

b.   Type—Network, Host, or Range.

c.   IP Address—An IPv4 or IPv6 address. If you select Range as the object type, the IP Address field changes to allow you to enter a Start Address and an End address.

    d. Netmask/Prefix Length—Enter the subnet mask or prefix length.

    e. Description—(Optional) The description of the network object (up to 200 characters in length).

Step 3 If the NAT section is hidden, click **NAT** to expand the section.

Step 4 Check the **Add Automatic Translation Rules** check box.

Step 5 From the Type drop-down list, choose **Static**.

Step 6 In the Translated Addr. field, do one of the following:

- Type the same IP address that you used for the real address.

- Click the browse button, and choose a network object with a matching IP address definition from the Browse Translated Addr dialog box.

- Click the browse button, and create a new network object with a matching IP address definition from the Browse Translated Addr dialog box.

Name	▲ ¹	IP Address	Netmask
⊟ IPv4 Network Objects			
— 🖳 A_10.1.1.1		10.1.1.1	255.255.255...
∎◻ DMZnetwork1		209.165.201.0	255.255.255...

Step 7   (Optional) Click **Advanced**, and configure the following options in the Advanced NAT Settings dialog box.

- Disable Proxy ARP on egress interface—Disables proxy ARP for incoming packets to the mapped IP addresses. See Mapped Addresses and Routing, page 5-22 for more information.

- (Routed mode; interface(s) specified) Lookup route table to locate egress interface—Determines the egress interface using a route lookup instead of using the interface specified in the NAT command. See Determining the Egress Interface, page 5-24 for more information.

- (Required for Transparent Firewall Mode) Interface:

    – Source Interface—Specifies the real interface where this NAT rule applies. By default, the rule applies to all interfaces.

    – Destination Interface—Specifies the mapped interface where this NAT rule applies. By default, the rule applies to all interfaces.

Do not configure any other options on this dialog box. When you are finished, click **OK**. You return to the Add/Edit Network Object dialog box.

Step 8   Click **OK**, and then **Apply**.

Because static rules are bidirectional (allowing initiation to and from the real host), the NAT Rules table show two rows for each static rule, one for each direction.

Configuring Per-Session PAT Rules

By default, all TCP PAT traffic and all UDP DNS traffic uses per-session PAT. To use multi-session PAT for traffic, you can configure per-session PAT rules: a permit rule uses per-session PAT, and a deny rule uses multi-session PAT. For more information about per-session vs. multi-session PAT, see Per-Session PAT vs. Multi-Session PAT (Version 9.0(1) and Later), page 5-11.

Defaults

By default, the following rules are installed:

- Permit TCP from any (IPv4 and IPv6) to any (IPv4 and IPv6)
- Permit UDP from any (IPv4 and IPv6) to domain

These rules do not appear in the rule table.

Note   You cannot remove these rules, and they always exist after any manually-created rules. Because rules are evaluated in order, you can override the default rules. For example, to completely negate these rules, you could add the following:

- Deny TCP from any (IPv4 and IPv6) to any (IPv4 and IPv6)
- Deny UDP from any (IPv4 and IPv6) to domain

Detailed Steps

Step 1   Choose **Configuration > Firewall > Advanced > Per-Session NAT Rules**, and click **Add > Add Per-Session NAT Rule**.

Step 2   Click **Permit** or **Deny**.
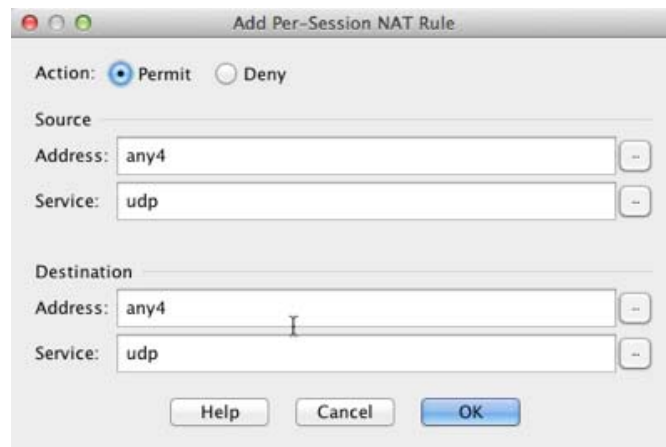
A permit rule uses per-session PAT; a deny rule uses multi-session PAT.

Step 3   Specify the Source Address either by typing an address or clicking the **...** button to choose an object.

Step 4   Specify the Source Service, UDP or TCP. You can optionally specify a source port, although normally you only specify the destination port. Either type in UDP/*port* or TCP/*port*, or click the **...** button to select a common value or object.

Step 5   Specify the Destination Address either by typing an address or clicking the **...** button to choose an object.

Step 6   Specify the Destination Service, UDP or TCP; this must match the source service. You can optionally specify a destination port. Either type in UDP/*port* or TCP/*port*, or click the **...** button to select a common value or object.

Step 7   Click **OK**.

Step 8   Click **Apply**.

Monitoring Network Object NAT

The Monitoring > Properties > Connection Graphs > Xlates pane lets you view the active Network Address Translations in a graphical format. You can choose up to four types of statistics to show in one graph window. You can open multiple graph windows at the same time.

Fields

- Available Graphs—Lists the components you can graph.

    - Xlate Utilization—Displays the ASA NAT utilization.

- Graph Window Title—Shows the graph window name to which you want to add a graph type. To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title.

- Add—Click to move the selected entries in the Available Graphs list to the Selected Graphs list.

- Remove—Click to remove the selected entry from the Selected Graphs list.

- Show Graphs—Click to display a new or updated graph window.

The Monitoring > Properties > Connection Graphs > Perfmon pane lets you view the performance information in a graphical format. You can choose up to four types of statistics to show in one graph window. You can open multiple graph windows at the same time.

Fields

- Available Graphs—Lists the components you can graph.
    - AAA Perfmon—Displays the ASA AAA performance information.
    - Inspection Perfmon—Displays the ASA inspection performance information.
    - Web Perfmon—Displays the ASA web performance information, including URL access and URL server requests.
    - Connections Perfmon—Displays the ASA connections performance information.
    - Xlate Perfmon—Displays the ASA NAT performance information.
- Graph Window Title—Shows the graph window name to which you want to add a graph type. To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title.
- Add—Click to move the selected entries in the Available Graphs list to the Selected Graphs list.
- Remove—Click to remove the selected statistic type from the Selected Graphs list.
- Show Graphs—Click to display a new or updated graph window.

Configuration Examples for Network Object NAT

This section includes the following configuration examples:

Providing Access to an Inside Web Server (Static 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. (See Figure 6-1).

Figure 6-1   Static NAT for an Inside Web Server
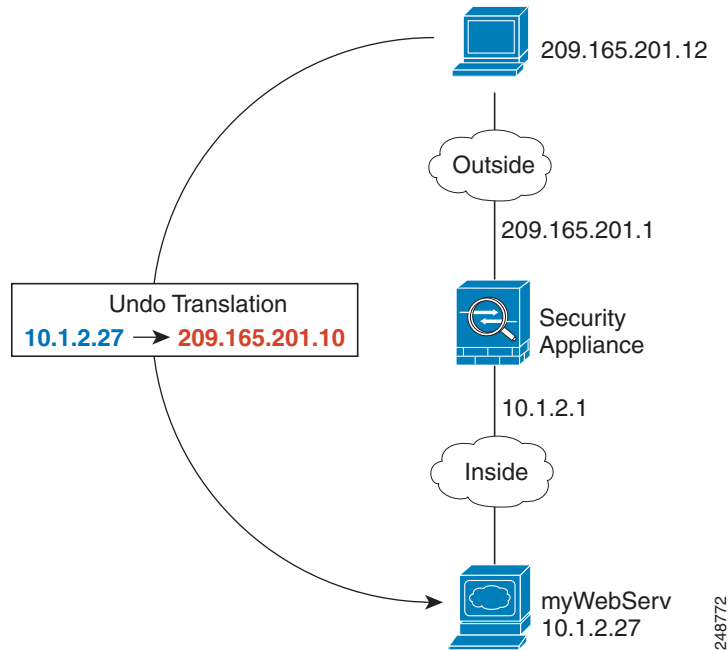
Step 1   Create a network object for the internal web server:
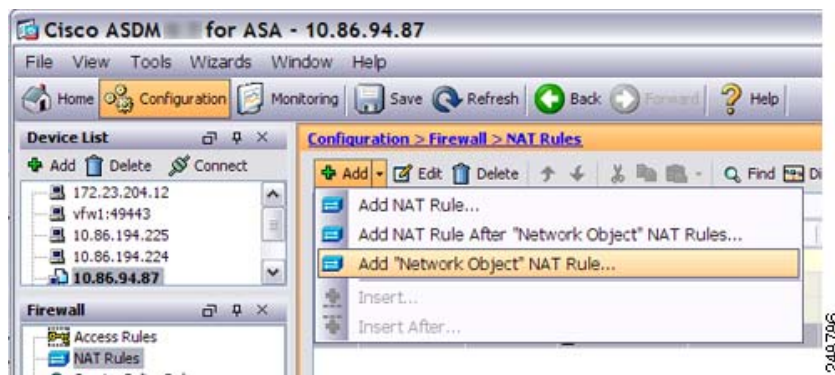
Step 2   Define the web server address:

Step 3   Configure static NAT for the object:
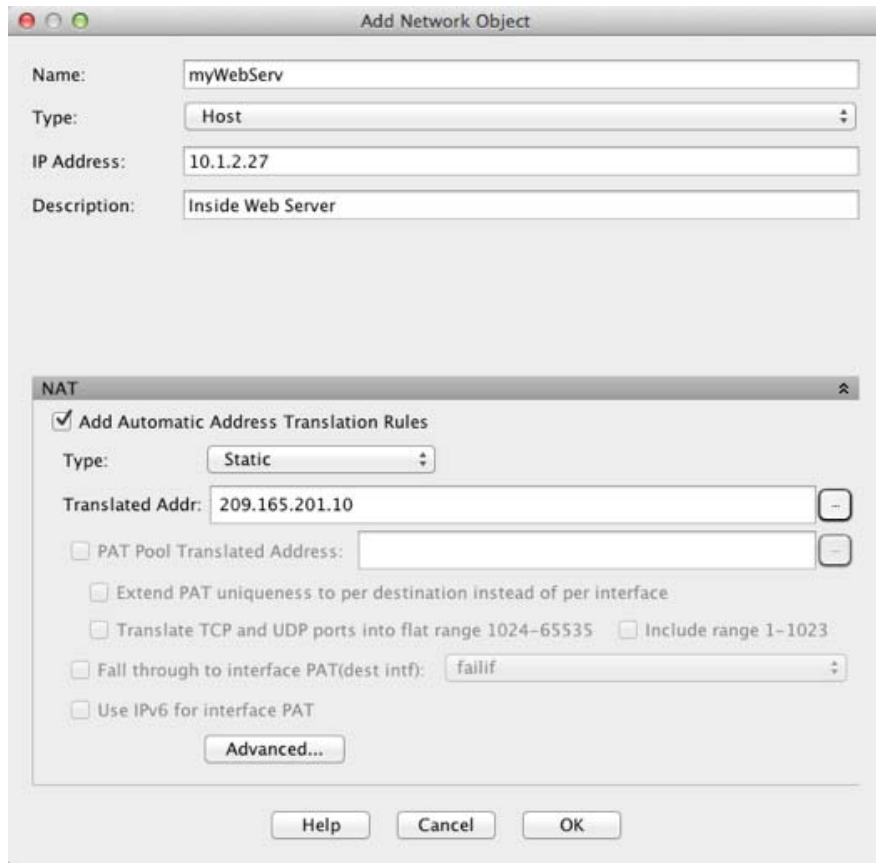
Step 4   Configure the real and mapped interfaces by clicking **Advanced**:
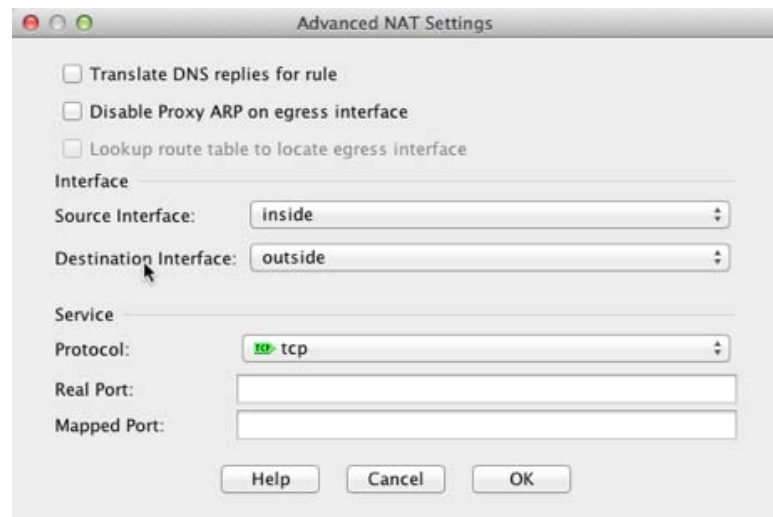
Step 5   Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

NAT for Inside Hosts (Dynamic NAT) and NAT for an Outside Web Server (Static NAT)

The following example configures dynamic NAT for inside users on a private network when they access the outside. Also, when inside users connect to an outside web server, that web server address is translated to an address that appears to be on the inside network. (See Figure 6-2).

Figure 6-2   *Dynamic NAT for Inside, Static NAT for Outside Web Server*
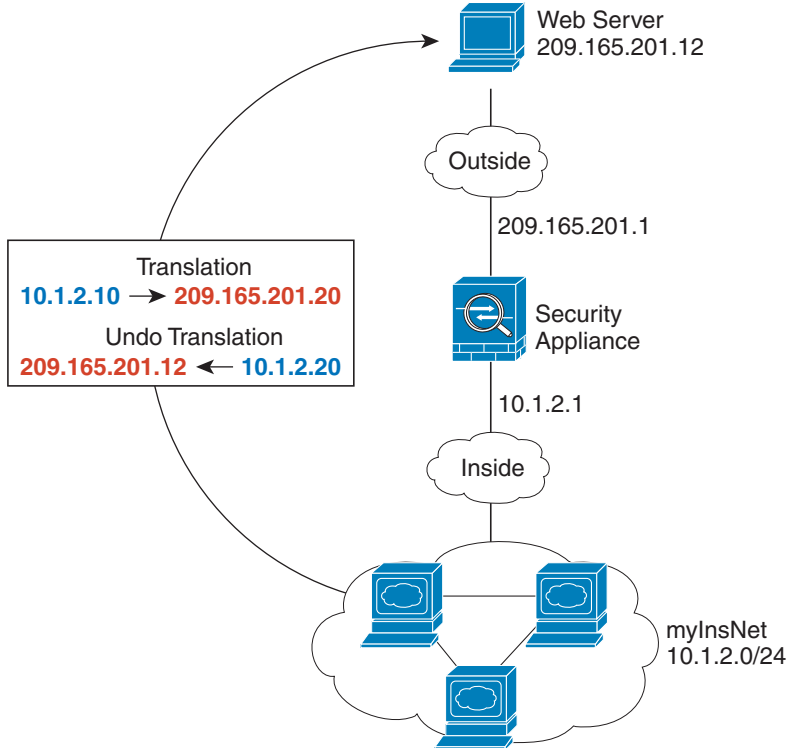
Step 1   Create a network object for the inside network:
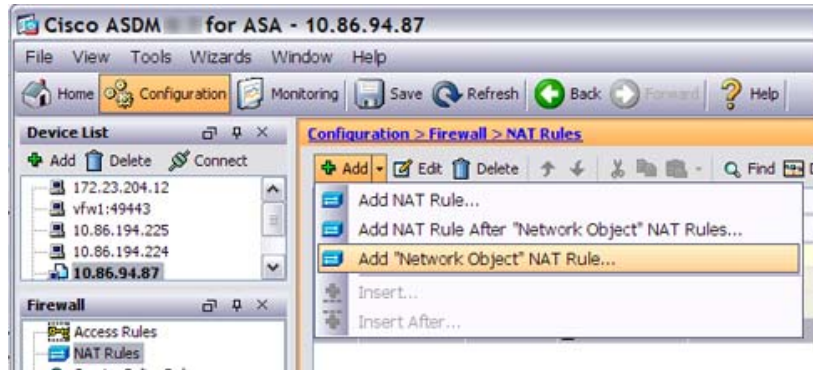
Step 2   Define the addresses for the inside network:

Step 3   Enable dynamic NAT for the inside network:
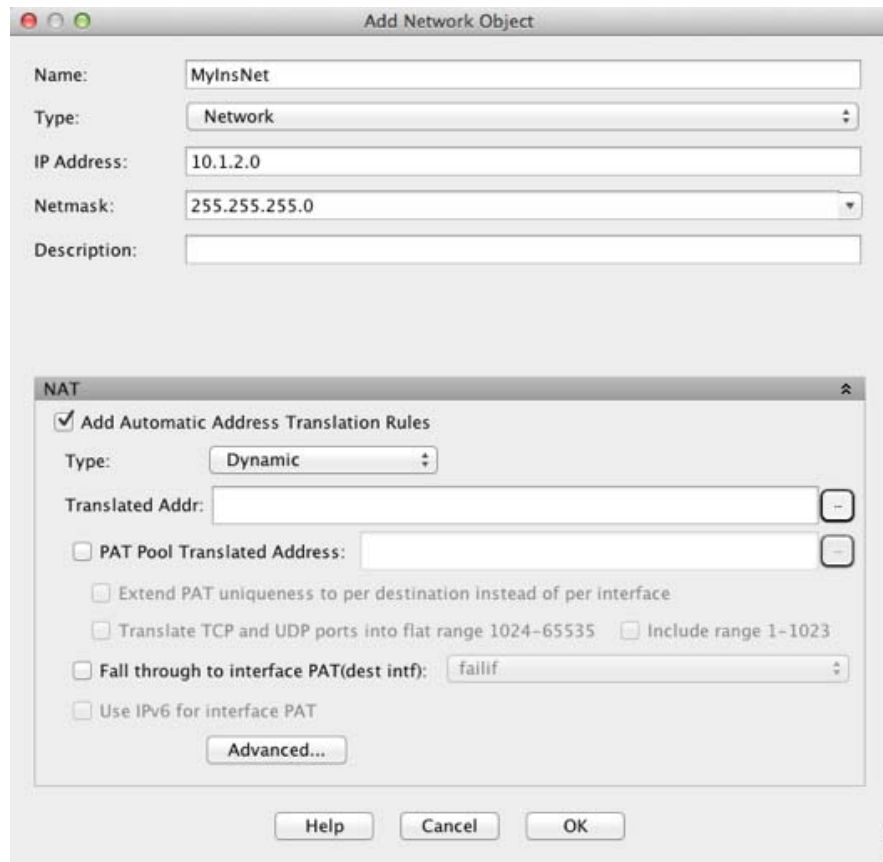
Step 4   For the Translated Addr field, add a new network object for the dynamic NAT pool to which you want to translate the inside addresses by clicking the browse button.

   a.   Add the new network object.

b. Define the NAT pool addresses, and click **OK**.

c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 5 Configure the real and mapped interfaces by clicking **Advanced**:
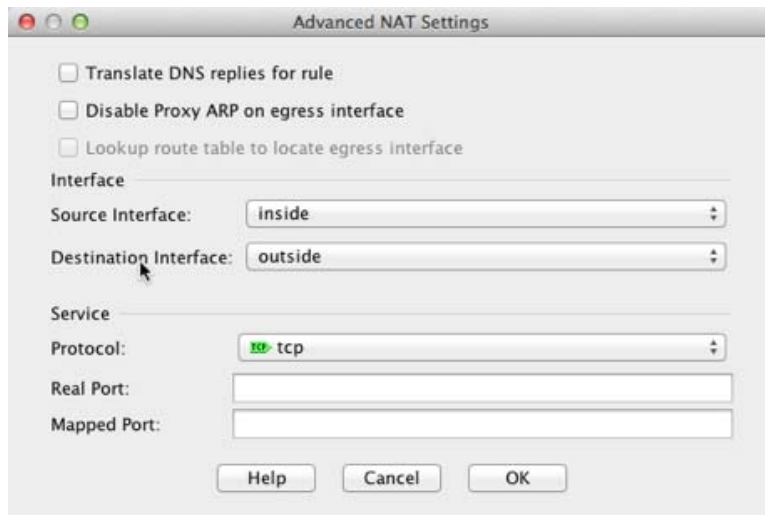
Step 6   Click **OK** to return to the Edit Network Object dialog box, click then click **OK** again to return to the NAT Rules table.

Step 7   Create a network object for the outside web server:
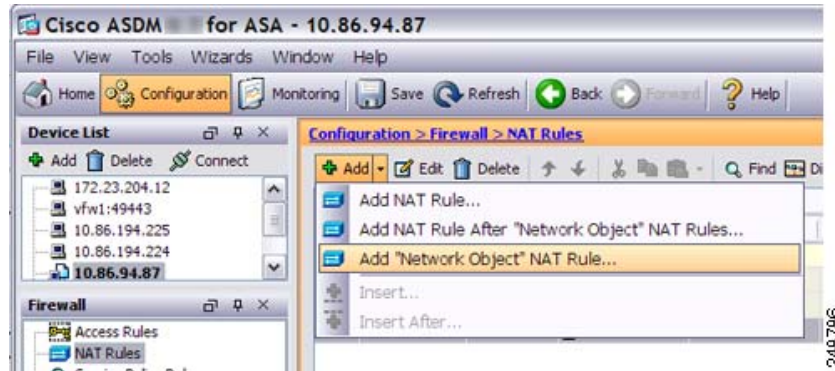
Step 8   Define the web server address:

Step 9   Configure static NAT for the web server:
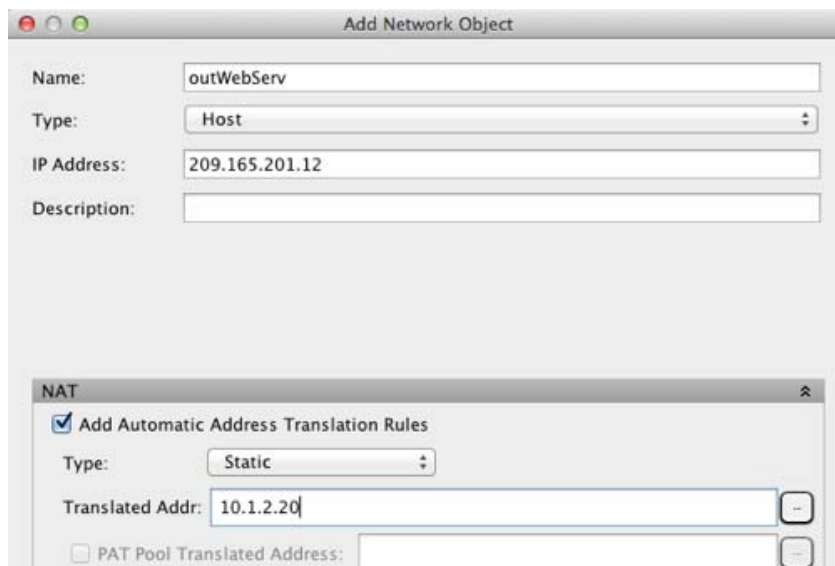
Step 10   Configure the real and mapped interfaces by clicking **Advanced**:
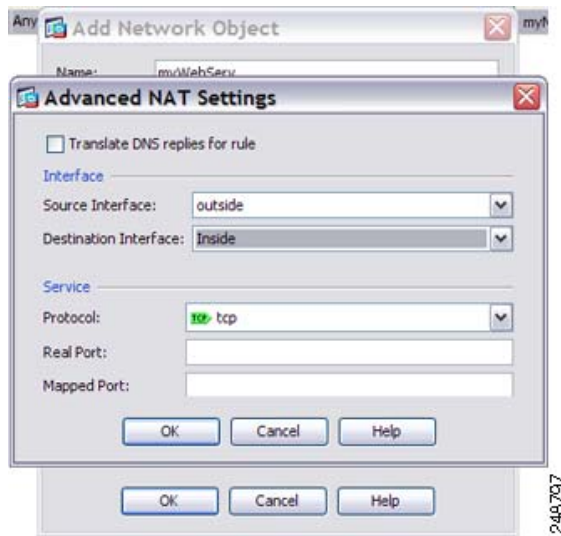
Step 11   Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

Inside Load Balancer with Multiple Mapped Addresses (Static NAT, One-to-Many)

The following example shows an inside load balancer that is translated to multiple IP addresses. When an outside host accesses one of the mapped IP addresses, it is untranslated to the single load balancer address. Depending on the URL requested, it redirects traffic to the correct web server. (See Figure 6-3).

Figure 6-3   *Static NAT with One-to-Many for an Inside Load Balancer*

Step 1   Create a network object for the load balancer:

Step 2   Define the load balancer address:

Step 3   Configure static NAT for the load balancer:
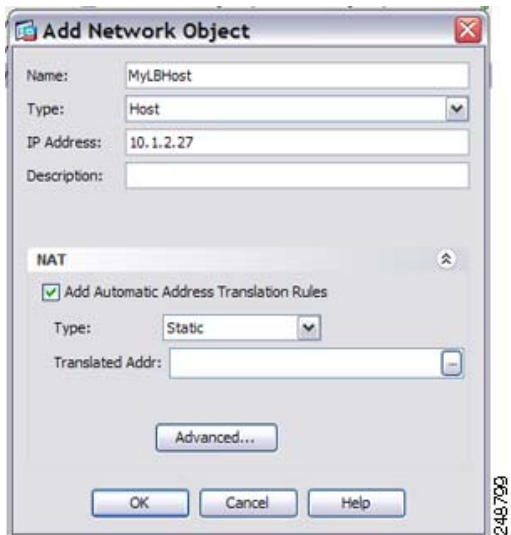
Step 4   For the Translated Addr field, add a new network object for the static NAT group of addresses to which you want to translate the load balancer address by clicking the browse button.

    a.   Add the new network object.

    b.   Define the static NAT group of addresses, and click **OK**.

c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 5   Configure the real and mapped interfaces by clicking **Advanced**:
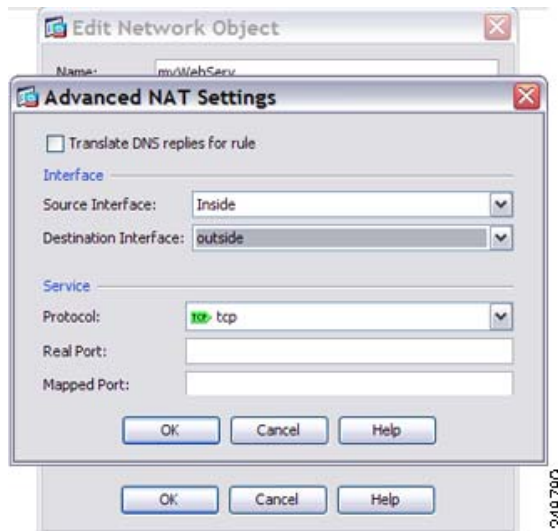
Step 6   Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

Single Address for FTP, HTTP, and SMTP (Static 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. (See Figure 6-4.)

Figure 6-4   Static NAT-with-Port-Translation

Step 1   Create a network object for the FTP server address:
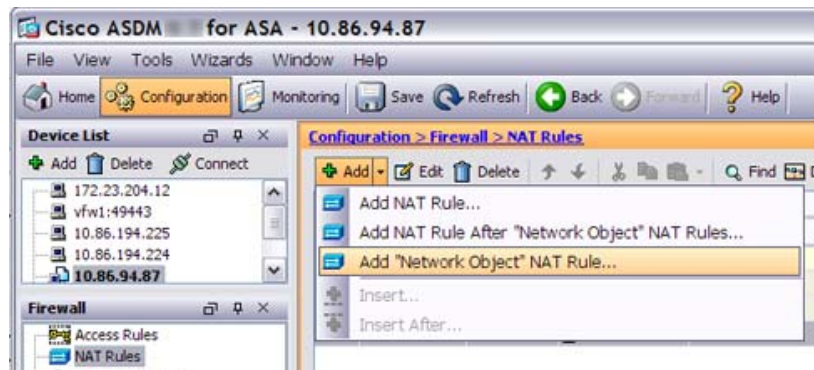
Step 2   Define the FTP server address, and configure static NAT with identity port translation for the FTP server:
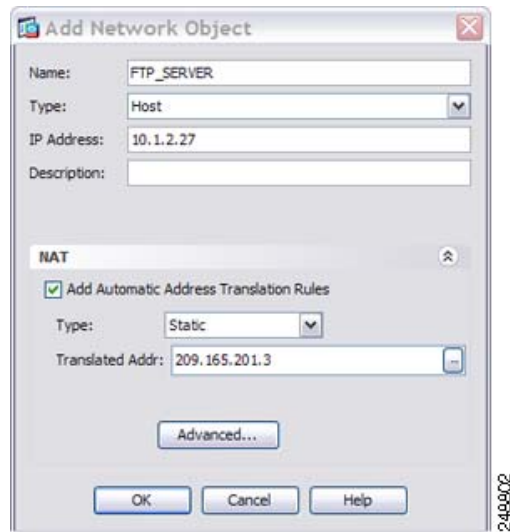
Step 3   Click **Advanced** to configure the real and mapped interfaces and port translation for FTP.
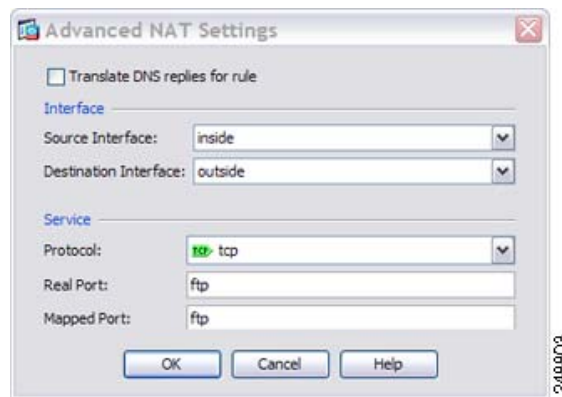
Step 4   Create a network object for the HTTP server address:

Step 5   Define the HTTP server address, and configure static NAT with identity port translation for the HTTP server:

Step 6   Click **Advanced** to configure the real and mapped interfaces and port translation for HTTP.

Step 7   Create a network object for the SMTP server address:

Step 8   Define the SMTP server address, and configure static NAT with identity port translation for the SMTP server:

Step 9   Click **Advanced** to configure the real and mapped interfaces and port translation for SMTP.

Step 10   Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

DNS Server on Mapped Interface, Web Server on Real Interface (Static NAT with DNS Modification)

For example, a DNS server is accessible from the outside interface. A server, ftp.cisco.com, is on the inside interface. You configure the ASA 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. (See Figure 6-5.) 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 ASA 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.

Figure 6-5   DNS Reply Modification
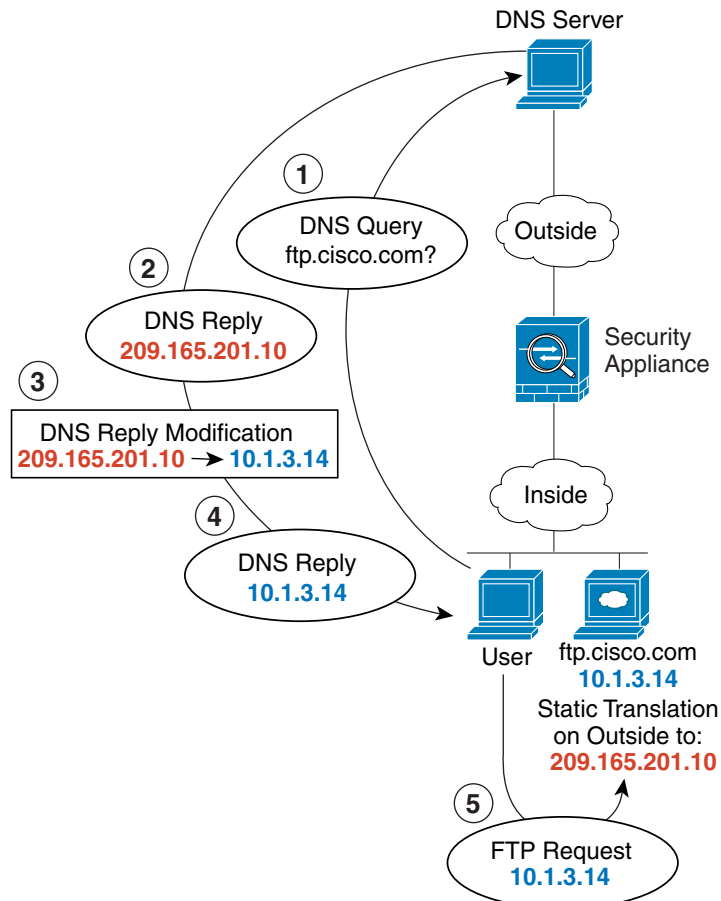
Step 1   Create a network object for the FTP server address:

Step 2   Define the FTP server address, and configure static NAT with DNS modification:
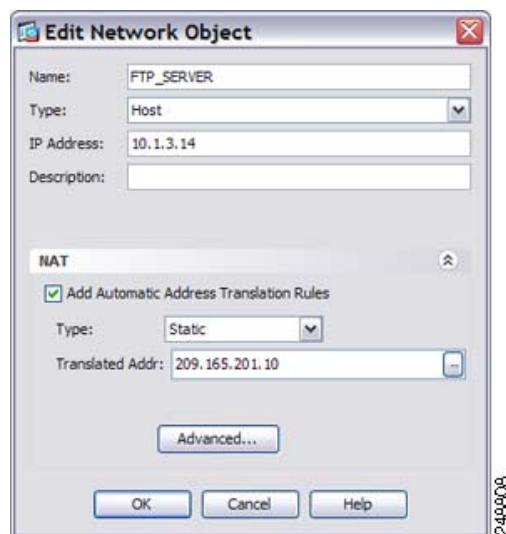
Step 3   Click **Advanced** to configure the real and mapped interfaces and DNS modification.
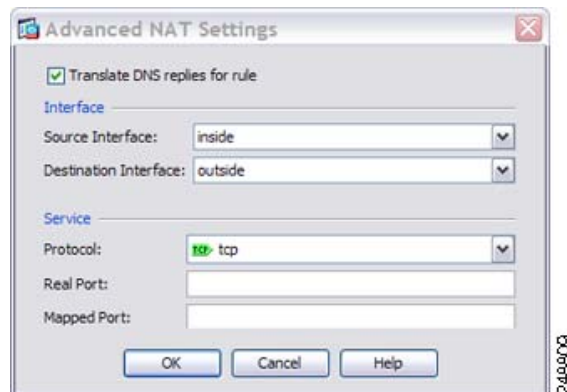
Step 4   Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

DNS Server and FTP Server on Mapped Interface, FTP Server is Translated (Static NAT with DNS Modification)

Figure 6-6 shows an FTP server and DNS server on the outside. The ASA 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, 209.165.201.10. 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.

Figure 6-6   DNS Reply Modification Using Outside NAT
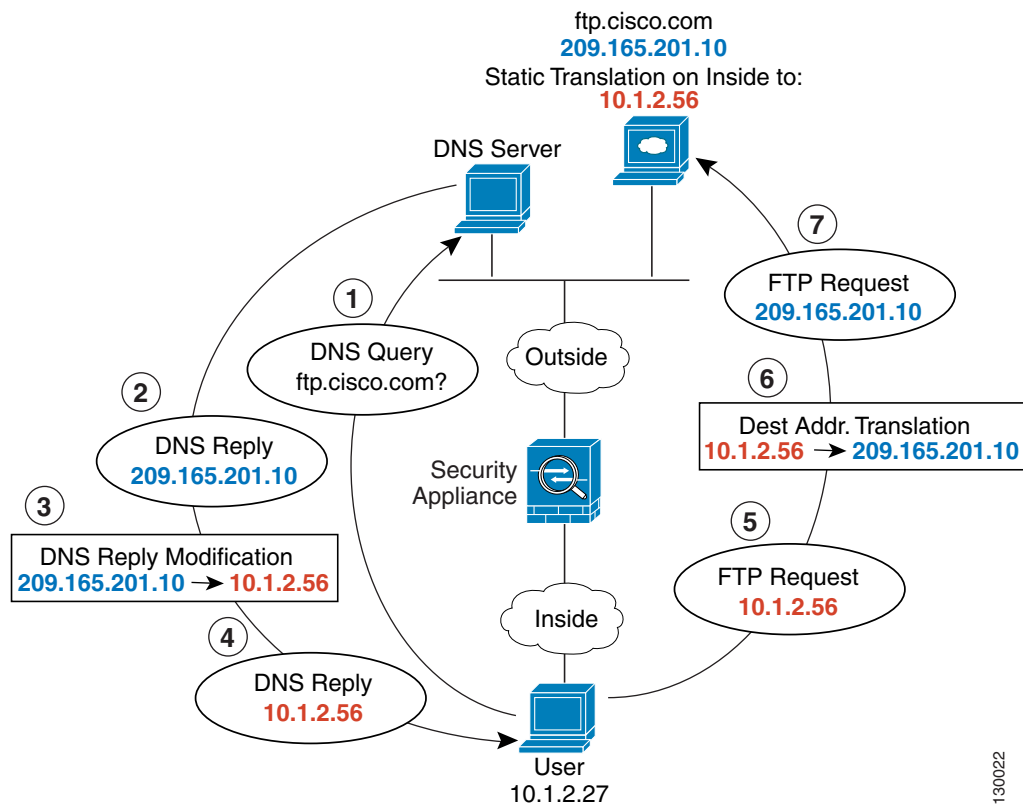
Step 1   Create a network object for the FTP server address:

Step 2   Define the FTP server address, and configure static NAT with DNS modification:
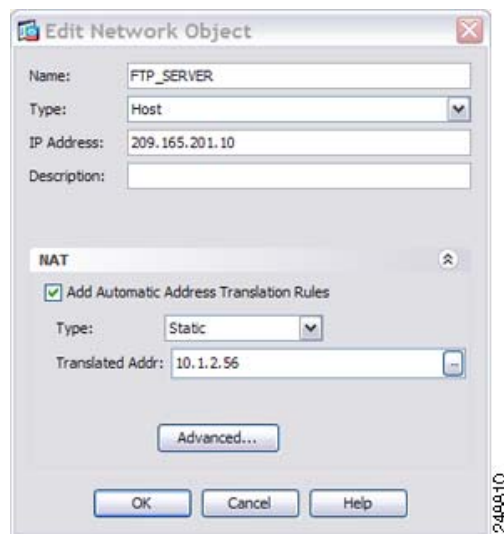
Step 3   Click **Advanced** to configure the real and mapped interfaces and DNS modification.
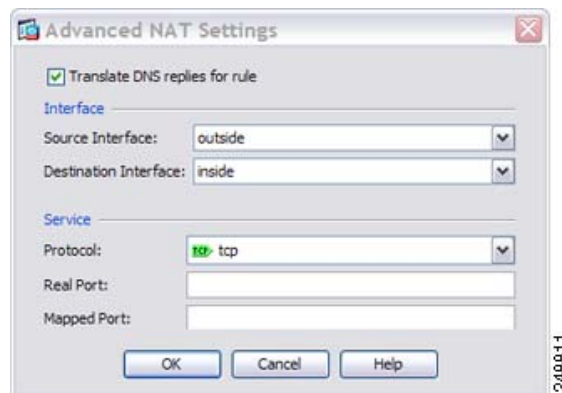
Step 4   Click **OK** to return to the Edit Network Object dialog box, click **OK** again, and then click **Apply**.

IPv4 DNS Server and FTP Server on Mapped Interface, IPv6 Host on Real Interface (Static NAT64 with DNS64 Modification)

Figure 6-6 shows an FTP server and DNS server on the outside IPv4 network. The ASA 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) 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.

Figure 6-7   DNS Reply Modification Using Outside NAT

Step 1   Configure static NAT with DNS modification for the FTP server.

    a.   Create a network object for the FTP server address.
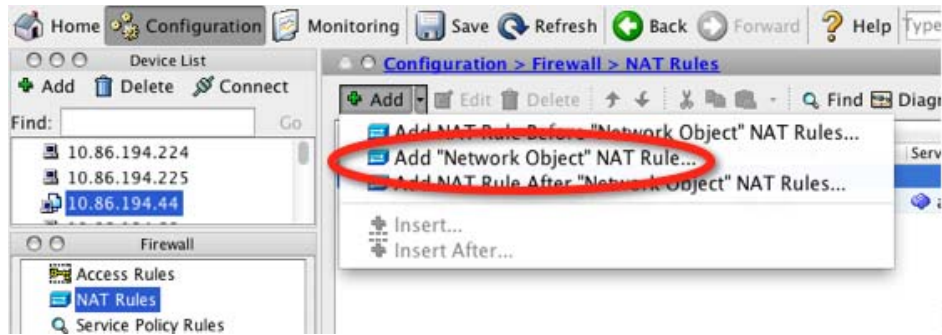
b. Define the FTP server address, and configure static NAT with DNS modification and, because this is a one-to-one translation, configure the one-to-one method for NAT46.
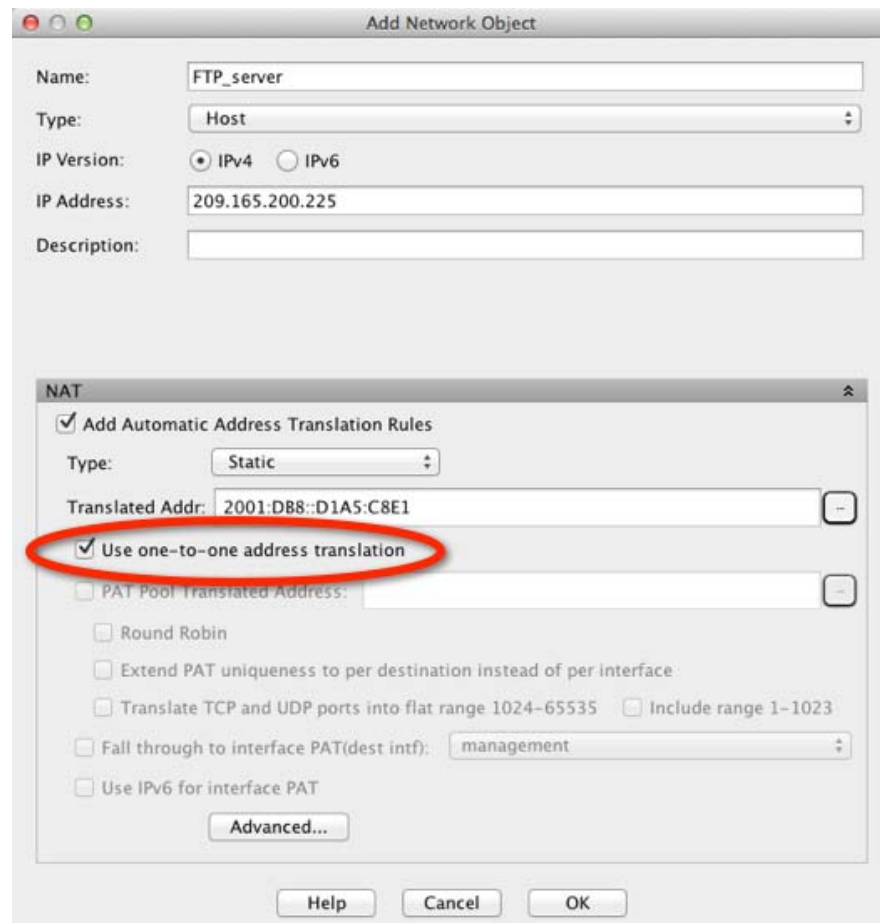
c. Click **Advanced** to configure the real and mapped interfaces and DNS modification.

d. Click **OK** to return to the Edit Network Object dialog box.

Step 2 Configure NAT for the DNS server.

a. Create a network object for the DNS server address.

b. Define the DNS server address, and configure static NAT using the one-to-one method.
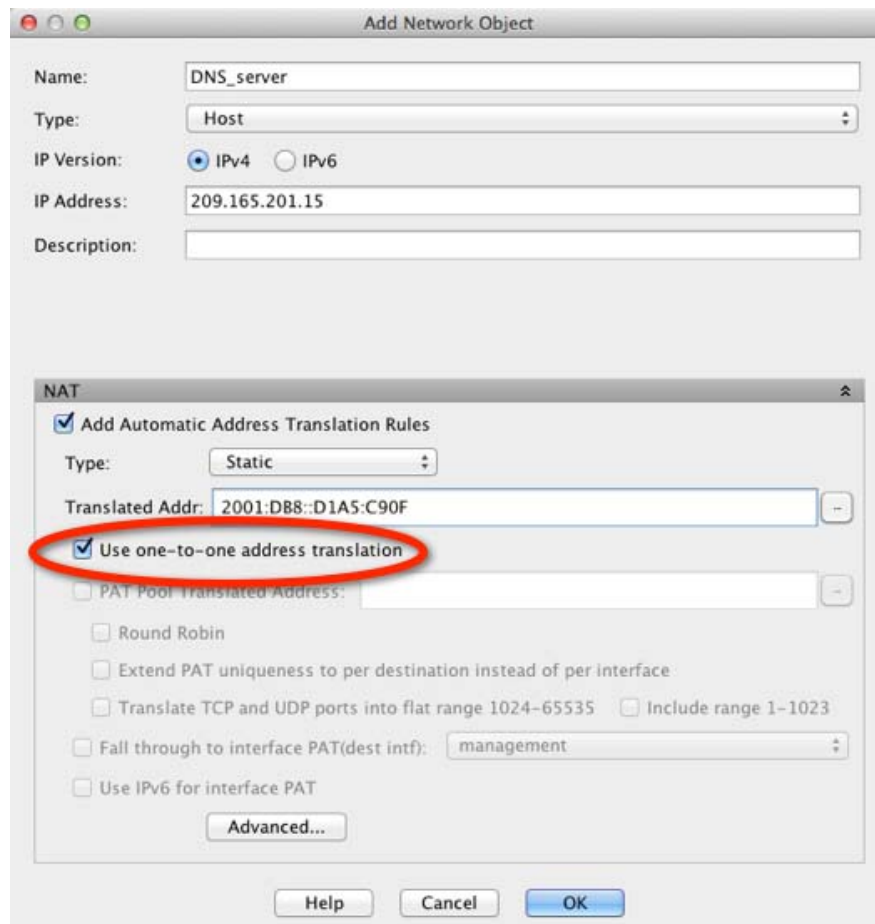
c. Click **Advanced** to configure the real and mapped interfaces.

d. Click **OK** to return to the Edit Network Object dialog box.

Step 3   Configure an IPv4 PAT pool for translating the inside IPv6 network.

Under NAT, uncheck the **Add Automatic Address Translation Rules** check box.

Step 4   Configure PAT for the inside IPv6 network.

a. Create a network object for the inside IPv6 network.

b. Define the IPv6 network address, and configure dynamic NAT using a PAT pool.

c. Next to the PAT Pool Translated Address field, click the **...** button to choose the PAT pool you created earlier, and click **OK**.
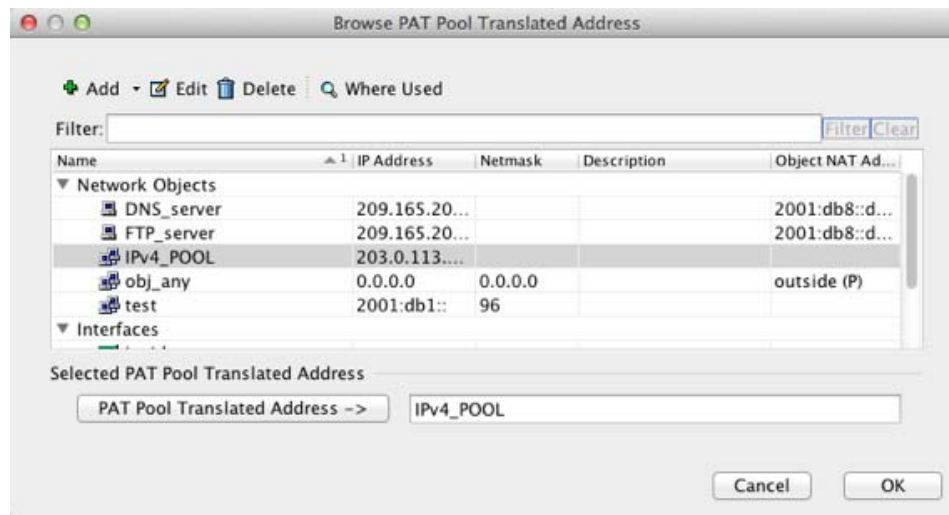
d. Click **Advanced** to configure the real and mapped interfaces.

     e. Click **OK** to return to the Edit Network Object dialog box.

Step 5   Click **OK**, and then click **Apply**.

Feature History for Network Object NAT

Table 6-1 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 6-1   Feature History for Network Object NAT

Feature Name	Platform Releases	Feature Information
Network Object NAT	8.3(1)	Configures NAT for a network object IP address(es).
		We introduced or modified the following screens: Configuration > Firewall > NAT Rules Configuration > Firewall > Objects > Network Objects/Groups
Identity NAT configurable proxy ARP and route lookup	8.4(2)/8.5(1)	In earlier releases for identity NAT, proxy ARP was disabled, and a route lookup was always used to determine the egress interface. You could not configure these settings. In 8.4(2) and later, the default behavior for identity NAT was changed to match the behavior of other static NAT configurations: proxy ARP is enabled, and the NAT configuration determines the egress interface (if specified) by default. You can leave these settings as is, or you can enable or disable them discretely. Note that you can now also disable proxy ARP for regular static NAT.
		When upgrading to 8.4(2) from 8.3(1), 8.3(2), and 8.4(1), all identity NAT configurations will now include the **no-proxy-arp** and **route-lookup** keywords, to maintain existing functionality.
		We modified the following screen: Configuration > Firewall > NAT Rules > Add/Edit Network Object > Advanced NAT Settings.

Table 6-1   *Feature History for Network Object NAT (continued)*

Feature Name	Platform Releases	Feature Information
PAT pool and round robin address assignment	8.4(2)/8.5(1)	You can now specify a pool of PAT addresses instead of a single address. You can also optionally enable round-robin assignment of PAT addresses instead of first using all ports on a PAT address before using the next address in the pool. These features help prevent a large number of connections from a single PAT address from appearing to be part of a DoS attack and makes configuration of large numbers of PAT addresses easy. We modified the following screens: Configuration > Firewall > NAT Rules > Add/Edit Network Object.
Round robin PAT pool allocation uses the same IP address for existing hosts	8.4(3)	When using a PAT pool with round robin allocation, if a host has an existing connection, then subsequent connections from that host will use the same PAT IP address if ports are available. We did not modify any screens. *This feature is not available in 8.5(1) or 8.6(1).*
Flat range of PAT ports for a PAT pool	8.4(3)	If available, the real source port number is used for the mapped port. However, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 0 to 511, 512 to 1023, and 1024 to 65535. Therefore, ports below 1024 have only a small PAT pool. If you have a lot of traffic that uses the lower port ranges, when using a PAT pool, you can now specify a flat range of ports to be used instead of the three unequal-sized tiers: either 1024 to 65535, or 1 to 65535. We modified the following screens: Configuration > Firewall > NAT Rules > Add/Edit Network Object. *This feature is not available in 8.5(1) or 8.6(1).*
Extended PAT for a PAT pool	8.4(3)	Each PAT IP address allows up to 65535 ports. If 65535 ports do not provide enough translations, you can now enable extended PAT for a PAT pool. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. We modified the following screens: Configuration > Firewall > NAT Rules > Add/Edit Network Object. *This feature is not available in 8.5(1) or 8.6(1).*

Table 6-1   ***Feature History for Network Object NAT (continued)***

Feature Name	Platform Releases	Feature Information
Automatic NAT rules to translate a VPN peer's local IP address back to the peer's real IP address	8.4(3)	In rare situations, you might want to use a VPN peer's real IP address on the inside network instead of an assigned local IP address. Normally with VPN, the peer is given an assigned local IP address to access the inside network. However, you might want to translate the local IP address back to the peer's real public IP address if, for example, your inside servers and network security is based on the peer's real IP address.
		You can enable this feature on one interface per tunnel group. Object NAT rules are dynamically added and deleted when the VPN session is established or disconnected. You can view the rules using the **show nat** command.
		Note Because of routing issues, we do not recommend using this feature unless you know you need this feature; contact Cisco TAC to confirm feature compatibility with your network. See the following limitations:
		• Only supports Cisco IPsec and AnyConnect Client.
		• Return traffic to the public IP addresses must be routed back to the ASA so the NAT policy and VPN policy can be applied.
		• Does not support load-balancing (because of routing issues).
		• Does not support roaming (public IP changing).
		ASDM does not support this command; enter the command using the Command Line Tool.
NAT support for IPv6	9.0(1)	NAT now supports IPv6 traffic, as well as translating between IPv4 and IPv6. Translating between IPv4 and IPv6 is not supported in transparent mode.
		We modified the following screen: Configuration > Firewall > Objects > Network Objects/Group.

Table 6-1   *Feature History for Network Object NAT (continued)*

Feature Name	Platform Releases	Feature Information
NAT support for reverse DNS lookups	9.0(1)	NAT now supports translation of the DNS PTR record for reverse DNS lookups when using IPv4 NAT, IPv6 NAT, and NAT64 with DNS inspection enabled for the NAT rule.
Per-session PAT	9.0(1)	The per-session PAT feature improves the scalability of PAT and, for clustering, allows each member unit to own PAT connections; multi-session PAT connections have to be forwarded to and owned by the master unit. At the end of a per-session PAT session, the ASA sends a reset and immediately removes the xlate. This reset causes the end node to immediately release the connection, avoiding the TIME_WAIT state. Multi-session PAT, on the other hand, uses the PAT timeout, by default 30 seconds. For "hit-and-run" traffic, such as HTTP or HTTPS, the per-session feature can dramatically increase the connection rate supported by one address. Without the per-session feature, the maximum connection rate for one address for an IP protocol is approximately 2000 per second. With the per-session feature, the connection rate for one address for an IP protocol is 65535/*average-lifetime*.

By default, all TCP traffic and UDP DNS traffic use a per-session PAT xlate. For traffic that requires multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT by creating a per-session deny rule.

We introduced the following screen: Configuration > Firewall > Advanced > Per-Session NAT Rules. |

Twice NAT (ASA 8.3 and Later)

Twice NAT lets you identify both the source and destination address in a single rule. This chapter shows you how to configure twice NAT and includes the following sections:

Note For detailed information about how NAT works, see Chapter 5, "Information About NAT (ASA 8.3 and Later)."

Information About Twice NAT

Twice 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 a source address should be translated to A when going to destination X, but be translated to B when going to destination Y, for example.

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 in the command, 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.

Twice NAT also lets you use service objects for static NAT-with-port-translation; network object NAT only accepts inline definition.

For detailed information about the differences between twice NAT and network object NAT, see How NAT is Implemented, page 5-15.

Twice NAT rules are added to section 1 of the NAT rules table, or if specified, section 3. For more information about NAT ordering, see NAT Rule Order, page 5-20.

Licensing Requirements for Twice NAT

Model	License Requirement
ASAv	Standard or Premium License.
All other models	Base License.

Prerequisites for Twice NAT

- For both the real and mapped addresses, configure network objects or network object groups. Network object groups are particularly useful for creating a mapped address pool with discontinuous IP address ranges or multiple hosts or subnets. To create a network object or group, see the general operations configuration guide.

- For static NAT-with-port-translation, configure TCP or UDP service objects. To create a service object, see the general operations configuration guide.

For specific guidelines for objects and groups, see the configuration section for the NAT type you want to configure. See also the Guidelines and Limitations, page 7-2 section.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

- Supported in routed and transparent firewall mode.

- In transparent mode, you must specify the real and mapped interfaces; you cannot use --Any--.

- In transparent mode, you cannot configure interface PAT, because the transparent mode interfaces do not have IP addresses. You also cannot use the management IP address as a mapped address.

- In transparent mode, translating between IPv4 and IPv6 networks is not supported. Translating between two IPv6 networks, or between two IPv4 networks is supported.

IPv6 Guidelines

- Supports IPv6.

- For routed mode, you can also translate between IPv4 and IPv6.

- For transparent mode, translating between IPv4 and IPv6 networks is not supported. Translating between two IPv6 networks, or between two IPv4 networks is supported.

- For transparent mode, a PAT pool is not supported for IPv6.

- 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.

Additional Guidelines

- You cannot configure FTP destination port translation when the source IP address is a subnet (or any other application that uses a secondary connection); the FTP data channel establishment does not succeed.

- If you change the NAT configuration, and you do not want to wait for existing translations to time out before the new NAT information is used, you can clear the translation table using the **clear xlate** command. However, clearing the translation table disconnects all current connections that use translations.

    > **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** command. 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.

- When using the **any** keyword in a NAT rule, the definition of "any" traffic (IPv4 vs. IPv6) depends on the rule. Before the ASA performs NAT on a packet, the packet must be IPv6-to-IPv6 or IPv4-to-IPv4; with this prerequisite, the ASA 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.

- Objects and object groups used in NAT cannot be undefined; they must include IP addresses.

- You can use the same objects in multiple 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), use the interface name instead of the IP address.

    - (Transparent mode) The management IP address.

    - (Dynamic NAT) The standby interface IP address when VPN is enabled.

    - Existing VPN pool addresses.

Default Settings

- By default, the rule is added to the end of section 1 of the NAT table.

- (Routed mode) The default real and mapped interface is Any, which applies the rule to all interfaces.

- (8.3(1), 8.3(2), and 8.4(1)) The default behavior for identity NAT has proxy ARP disabled. You cannot configure this setting. (8.4(2) and later) The default behavior for identity NAT has proxy ARP enabled, matching other static NAT rules. You can disable proxy ARP if desired.

- If you specify an optional interface, then the ASA uses the NAT configuration to determine the egress interface. (8.3(1) through 8.4(1)) The only exception is for identity NAT, which always uses a route lookup, regardless of the NAT configuration. (8.4(2) and later) For identity NAT, the default behavior is to use the NAT configuration, but you have the option to always use a route lookup instead.

Configuring Twice NAT

This section describes how to configure twice NAT. This section includes the following topics:

Configuring Dynamic NAT or Dynamic PAT Using a PAT Pool

This section describes how to configure twice NAT for dynamic NAT or for dynamic PAT using a PAT pool. For more information, see Dynamic NAT, page 5-8 or the Dynamic PAT, page 5-10.

Guidelines

For a PAT pool:

- If available, the real source port number is used for the mapped port. However, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 0 to 511, 512 to 1023, and 1024 to 65535. Therefore, ports below 1024 have only a small PAT pool that can be used. (8.4(3) and later, not including 8.5(1) or 8.6(1)) If you have a lot of traffic that uses the lower port ranges, you can now specify for a PAT pool a flat range of ports to be used instead of the three unequal-sized tiers: either 1024 to 65535, or 1 to 65535.

- (8.4(3) and later, not including 8.5(1) or 8.6(1)) If you use the same PAT pool object in two separate rules, then be sure to specify the same options for each rule. For example, if one rule specifies extended PAT and a flat range, then the other rule must also specify extended PAT and a flat range.

For extended PAT for a PAT pool (8.4(3) and later, not including 8.5(1) or 8.6(1)):

- Many application inspections do not support extended PAT. See Default Settings and NAT Limitations, page 8-4 in Chapter 8, "Getting Started with Application Layer Protocol Inspection," for a complete list of unsupported inspections.

- If you enable extended PAT for a dynamic PAT rule, then you cannot also use an address in the PAT pool as the PAT address in a separate static NAT with port translation rule. For example, if the PAT pool includes 10.1.1.1, then you cannot create a static NAT-with-port-translation rule using 10.1.1.1 as the PAT address.

- Extended PAT can consume a large amount of memory because NAT pools are created for each unique destination, which in turn uses up memory. This may lead to memory exhaustion quickly even with less number of connections.

- If you use a PAT pool and specify an interface for fallback, you cannot specify extended PAT.

- For VoIP deployments that use ICE or TURN, do not use extended PAT. ICE and TURN rely on the PAT binding to be the same for all destinations.

For round robin for a PAT pool:

- (8.4(3) and later, not including 8.5(1) or 8.6(1)) If a host has an existing connection, then subsequent connections from that host will use the same PAT IP address if ports are available. **Note**: This "stickiness" does not survive a failover. If the ASA fails over, then subsequent connections from a host may not use the initial IP address.

- (8.4(2), 8.5(1), and 8.6(1)) If a host has an existing connection, then subsequent connections from that host will likely use *different* PAT addresses for each connection because of the round robin allocation. In this case, you may have problems when accessing two websites that exchange information about the host, for example an e-commerce site and a payment site. When these sites see two different IP addresses for what is supposed to be a single host, the transaction may fail.

- Round robin, especially when combined with extended PAT, can consume a large amount of memory. Because NAT pools are created for every mapped protocol/IP address/port range, round robin results in a large number of concurrent NAT pools, which use memory. Extended PAT results in an even larger number of concurrent NAT pools.

Detailed Steps

To configure dynamic NAT, perform the following steps:

Step 1   Choose **Configuration > Firewall > NAT Rules**, and then click **Add**.

If you want to add this rule to section 3 after the network object rules, then click the down arrow next to Add, and choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.

Step 2 Set the source and destination interfaces.

By default in routed mode, both interfaces are set to --Any--. In transparent firewall mode, you must set specific interfaces.

a. From the Match Criteria: Original Packet > Source Interface drop-down list, choose the source interface.

b. From the Match Criteria: Original Packet > Destination Interface drop-down list, choose the destination interface.

Step 3 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following figure for an example of the original packet vs. the translated packet.

Source

Inside

Destination

Outside

NAT

Real: 10.1.2.2
Mapped: 192.168.2.2

Real: 192.168.1.1
Mapped: 10.1.1.1

10.1.2.2 ---> 10.1.1.1

Original Packet

192.168.2.2 ---> 192.168.1.1

Translated Packet

a. For the Match Criteria: Original Packet > Source Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**.

Name	▲ 1	IP Address	Netmask
⊟ IPv4 Network Objects			
🖳 A_10.1.1.1		10.1.1.1	255.255.255...
🖳 DMZnetwork1		209.165.201.0	255.255.255...

b. (Optional) For the Match Criteria: Original Packet > Destination Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Destination Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see Main Differences Between Network Object NAT and Twice NAT, page 5-15.

Step 4   (Optional) Identify the original packet port (the *mapped destination port*). For the Match Criteria: Original Packet > Service, click the browse button and choose an existing TCP or UDP service object or create a new object from the Browse Original Service dialog box.

Dynamic NAT does not support port translation. However, because the destination translation is always static, you can perform port translation for the destination port. A service object can contain both a source and destination port, but only the destination port is used in this case. If you specify the source port, it will be ignored. 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. The "not equal" (**!=**) operator is not supported.

Step 5 Choose **Dynamic** from the Match Criteria: Translated Packet > Source NAT Type drop-down list.

This setting only applies to the source address; the destination translation is always static.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*). You can translate between IPv4 and IPv6 if desired. See the following figure for an example of the original packet vs. the translated packet.

a.   You can perform either dynamic NAT or Dynamic PAT using a PAT pool:

•   Dynamic NAT—For the Match Criteria: Translated Packet > Source Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Translated Source Address dialog box.

For dynamic NAT, you typically configure a larger group of source addresses to be mapped to a smaller group.

> ✎
>
> **Note** The object or group cannot contain a subnet.

- Dynamic PAT using a PAT pool—.To configure a PAT pool, check the **PAT Pool Translated Address** check box, then click the browse button and choose an existing network object or group or create a new object or group from the Browse Translated PAT Pool Address dialog box. **Note**: Leave the Source Address field empty.

Action: Translated Packet

Source NAT Type: Dynamic   ▲▼

Source Address: -- Original --   ..   Destination Address:

☑ PAT Pool Translated Address: _____ ..   Service:

  ☐ Round Robin

  ☐ Extend PAT uniqueness to per destination instead of per interface

  ☐ Translate TCP and UDP ports into flat range 1024-65535   ☐ Include range 1-1023

☐ Fall through to interface PAT

> ✎
>
> **Note** The object or group cannot contain a subnet.

(Optional) For a PAT pool, configure the following options:

- To assign addresses/ports in a round-robin fashion, check the **Round Robin** check box. Without round-robin, by default, all ports for a PAT address will be allocated before the next PAT address is used. The round-robin method assigns an address/port from each PAT address in the pool before returning to use the first address again, and then the second address, and so on.

- (8.4(3) and later, not including 8.5(1) or 8.6(1)) Check the **Extend PAT uniqueness to per destination instead of per interface** check box to use extended PAT. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80.

- (8.4(3) and later, not including 8.5(1) or 8.6(1)) Check the **Translate TCP or UDP ports into flat range (1024-65535)** check box to use the 1024 to 65535 port range as a single flat range when allocating ports. When choosing the mapped port number for a translation, the ASA uses the real source port number if it is available. However, without this option, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 1 to 511, 512 to 1023, and 1024 to 65535. To avoid running out of ports at the low ranges, configure this setting. To use the entire range of 1 to 65535, also check the **Include range 1 to 1023** check box.

  b. (Optional, Routed Mode Only) To use the interface IP address as a backup method if the other mapped source addresses are already allocated, check the **Fall through to interface PAT** check box. To use the IPv6 interface address, also check the **Use IPv6 for interface PAT** check box.

  The destination interface IP address is used. This option is only available if you configure a specific Destination Interface.

c. For the Match Criteria: Translated Packet > Destination Address, click the browse button and choose an existing network object, group, or interface or create a new object or group from the Browse Translated Destination Address dialog box.

For identity NAT for the destination address, simply use the same object or group for both the real and mapped addresses.

If you want to translate the destination address, then the static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired. For more information, see Static NAT, page 5-3. See Guidelines and Limitations, page 7-2 for information about disallowed mapped IP addresses.

For static interface NAT with port translation only, choose an interface from the Browse dialog box. Be sure to also configure a service translation (see Step 7). For this option, you must configure a specific interface for the Source Interface in Step 2. See Static Interface NAT with Port Translation, page 5-6 for more information.

Step 7   (Optional) Identify the translated packet port (the *real destination port*). For the Match Criteria: Translated Packet > Service, click the browse button and choose an existing TCP or UDP service object or create a new object from the Browse Translated Service dialog box.

Dynamic NAT does not support port translation. However, because the destination translation is always static, you can perform port translation for the destination port. A service object can contain both a source and destination port, but only the destination port is used in this case. If you specify the source port, it will be ignored. 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. The "not equal" (**!=**) operator is not supported.

Step 8   (Optional) Configure NAT options in the Options area.

    a. Enable rule —Enables this NAT rule. The rule is enabled by default.

    b. (For a source-only rule) Translate DNS replies that match this rule—Rewrites the DNS A record in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). You cannot configure DNS modification if you configure a destination address. See DNS and NAT, page 5-31 for more information.

    c. Description—Adds a description about the rule up to 200 characters in length.

Step 9   Click **OK**.

Configuring Dynamic PAT (Hide)

This section describes how to configure twice NAT for dynamic PAT (hide). For dynamic PAT using a PAT pool, see Configuring Dynamic NAT or Dynamic PAT Using a PAT Pool, page 7-4 instead of using this section. For more information, see Dynamic PAT, page 5-10.

Detailed Steps

To configure dynamic PAT, perform the following steps:

Step 1 Choose **Configuration > Firewall > NAT Rules**, and then click **Add**.

If you want to add this rule to section 3 after the network object rules, then click the down arrow next to Add, and choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.

Step 2   Set the source and destination interfaces.

By default in routed mode, both interfaces are set to --Any--. In transparent firewall mode, you must set specific interfaces.

a. From the Match Criteria: Original Packet > Source Interface drop-down list, choose the source interface.

b. From the Match Criteria: Original Packet > Destination Interface drop-down list, choose the destination interface.

Step 3   Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following figure for an example of the original packet vs. the translated packet.

a. For the Match Criteria: Original Packet > Source Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**.

b. (Optional) For the Match Criteria: Original Packet > Destination Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Destination Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see Main Differences Between Network Object NAT and Twice NAT, page 5-15.

Step 4 (Optional) Identify the original packet port (the *mapped destination port*). For the Match Criteria: Original Packet > Service, click the browse button and choose an existing TCP or UDP service object or create a new object from the Browse Original Service dialog box.

Dynamic PAT does not support additional port translation. However, because the destination translation is always static, you can perform port translation for the destination port. A service object can contain both a source and destination port, but only the destination port is used in this case. If you specify the source port, it will be ignored. 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. The "not equal" (**!=**) operator is not supported.

Step 5 Choose **Dynamic PAT (Hide)** from the Match Criteria: Translated Packet > Source NAT Type drop-down list.

This setting only applies to the source address; the destination translation is always static.

> ✎
>
> **Note** To configure dynamic PAT using a PAT pool, choose **Dynamic** instead of Dynamic PAT (Hide), see Configuring Dynamic NAT or Dynamic PAT Using a PAT Pool, page 7-4.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*). You can translate between IPv4 and IPv6 if desired. See the following figure for an example of the original packet vs. the translated packet.

a. For the Match Criteria: Translated Packet > Source Address, click the browse button and choose an existing network object or interface or create a new object from the Browse Translated Source Address dialog box.

If you want to use the IPv6 address of the interface, check the **Use IPv6 for interface PAT** check box.
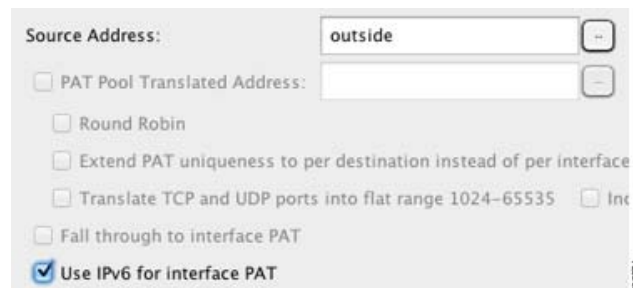
b. For the Match Criteria: Translated Packet > Destination Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Translated Destination Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

For identity NAT for the destination address, simply use the same object or group for both the real and mapped addresses.

If you want to translate the destination address, then the static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired. For more information, see Static NAT, page 5-3. See Guidelines and Limitations, page 7-2 for information about disallowed mapped IP addresses.

For static interface NAT with port translation only, choose an interface from the Browse dialog box. Be sure to also configure a service translation (see Step 7). For this option, you must configure a specific interface for the Source Interface in Step 2. See Static Interface NAT with Port Translation, page 5-6 for more information.

Step 7 (Optional) Identify the translated packet port (the *real destination port*). For the Match Criteria: Translated Packet > Service, click the browse button and choose an existing TCP or UDP service object from the Browse Translated Service dialog box.

You can also create a new service object from the Browse Translated Service dialog box and use this object as the mapped destination port.

Dynamic PAT does not support additional port translation. However, because the destination translation is always static, you can perform port translation for the destination port. A service object can contain both a source and destination port, but only the destination port is used in this case. If you specify the source port, it will be ignored. 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. The "not equal" (**!=**) operator is not supported.

Step 8   (Optional) Configure NAT options in the Options area.

a.   Enable rule —Enables this NAT rule. The rule is enabled by default.

b.   (For a source-only rule) Translate DNS replies that match this rule—Rewrites the DNS A record in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). You cannot configure DNS modification if you configure a destination address. See DNS and NAT, page 5-31 for more information.

c.   Description—Adds a description about the rule up to 200 characters in length.

Step 9   Click **OK**.

Configuring Static NAT or Static NAT-with-Port-Translation

This section describes how to configure a static NAT rule using twice NAT. For more information about static NAT, see Static NAT, page 5-3.

Detailed Steps

To configure static NAT, perform the following steps:

Step 1   Choose **Configuration > Firewall > NAT Rules**, and then click **Add**.

If you want to add this rule to section 3 after the network object rules, then click the down arrow next to Add, and choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.

Step 2 Set the source and destination interfaces.

By default in routed mode, both interfaces are set to --Any--. In transparent firewall mode, you must set specific interfaces.

   a. From the Match Criteria: Original Packet > Source Interface drop-down list, choose the source interface.

   b. From the Match Criteria: Original Packet > Destination Interface drop-down list, choose the destination interface.

Step 3 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). See the following figure for an example of the original packet vs. the translated packet.

Source **Destination**

Inside   NAT   Outside

Real: 10.1.2.2
Mapped: 192.168.2.2

Real: 192.168.1.1
Mapped: 10.1.1.1

10.1.2.2 ---> 10.1.1.1

192.168.2.2 ---> 192.168.1.1

Original Packet

Translated Packet

a. For the Match Criteria: Original Packet > Source Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**, but do not use this option except for identity NAT. See Configuring Identity NAT, page 7-24 for more information.

Name	▲ ¹	IP Address	Netmask
⊟ IPv4 Network Objects			
A_10.1.1.1		10.1.1.1	255.255.255...
DMZnetwork1		209.165.201.0	255.255.255...

b. (Optional) For the Match Criteria: Original Packet > Destination Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Destination Address dialog box.

Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see Main Differences Between Network Object NAT and Twice NAT, page 5-15.

Step 4 (Optional) Identify the original packet source or destination port (the *real source port* **or** the *mapped destination port*). For the Match Criteria: Original Packet > Service, click the browse button and choose an existing TCP or UDP service object or create a new object from the Browse Original Service dialog box.

A service object can contain both a source and destination port. You should specify *either* the source *or* the destination port for both the real and mapped service objects. You should only specify *both* the source and destination ports if your application uses a fixed source port (such as some DNS servers); but fixed source ports are rare. In the rare case where you specify both the source and destination ports in the object, the original packet service object contains the real source port/mapped destination port; the translated packet service object contains the mapped source port/real 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. The "not equal" (**!=**) operator is not supported.

Step 5   Choose **Static** from the Match Criteria: Translated Packet > Source NAT Type drop-down list. Static is
the default setting.

This setting only applies to the source address; the destination translation is always static.

Step 6   Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear
on the destination interface network (the *mapped source address* and the *real destination address*). You
can translate between IPv4 and IPv6 if desired. See the following figure for an example of the original
packet vs. the translated packet.

a.   For the Match Criteria: Translated Packet > Source Address, click the browse button and choose an
existing network object or group or create a new object or group from the Browse Translated Source
Address dialog box.

For static NAT, the mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired.

For static interface NAT with port translation, you can specify the interface instead of a network object/group for the mapped address. If you want to use the IPv6 address of the interface, check the **Use IPv6 for interface PAT** check box.

For more information, see Static Interface NAT with Port Translation, page 5-6. See Guidelines and Limitations, page 7-2 for information about disallowed mapped IP addresses.

b. For the Match Criteria: Translated Packet > Destination Address, click the browse button and choose an existing network object, group, or interface or create a new object or group from the Browse Translated Destination Address dialog box.

For static NAT, the mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired.

For static interface NAT with port translation, you can specify the interface instead of a network object/group for the mapped address. For more information, see Static Interface NAT with Port Translation, page 5-6. See Guidelines and Limitations, page 7-2 for information about disallowed mapped IP addresses.

Step 7   (Optional) Identify the translated packet source or destination port (the *mapped source port* **or** the *real destination port*). For the Match Criteria: Translated Packet > Service, click the browse button and choose an existing TCP or UDP service object or create a new object from the Browse Translated Service dialog box.

A service object can contain both a source and destination port. You should specify *either* the source *or* the destination port for both real and mapped service objects. You should only specify *both* the source and destination ports if your application uses a fixed source port (such as some DNS servers); but fixed source ports are rare. In the rare case where you specify both the source and destination ports in the object, the original packet service object contains the real source port/mapped destination port; the translated packet service object contains the mapped source port/real 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. The "not equal" (**!=**) operator is not supported.

Step 8   (Optional) For NAT46, check the **Use one-to-one address translation** check box. For NAT46, specify one-to-one to translate the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used. For a one-to-one translation, you must use this keyword.

Step 9   (Optional) Configure NAT options in the Options area.

   a.   Enable rule —Enables this NAT rule. The rule is enabled by default.

   b.   (For a source-only rule) Translate DNS replies that match this rule—Rewrites the DNS A record in DNS replies. Be sure DNS inspection is enabled (it is enabled by default). You cannot configure DNS modification if you configure a destination address. See DNS and NAT, page 5-31 for more information.

   c.   Disable Proxy ARP on egress interface—Disables proxy ARP for incoming packets to the mapped IP addresses. See Mapped Addresses and Routing, page 5-22 for more information.

   d.   Direction—To make the rule unidirectional, choose **Unidirectional**. The default is Both. Making the rule unidirectional prevents traffic from initiating connections to the real addresses.

   e.   Description—Adds a description about the rule up to 200 characters in length.

Step 10   Click **OK**.

Configuring Identity NAT

This section describes how to configure an identity NAT rule using twice NAT. For more information about identity NAT, see Identity NAT, page 5-12.

Detailed Steps

To configure identity NAT, perform the following steps:

Step 1   Choose **Configuration > Firewall > NAT Rules**, and then click **Add**.

If you want to add this rule to section 3 after the network object rules, then click the down arrow next to Add, and choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.

Step 2   Set the source and destination interfaces.

By default in routed mode, both interfaces are set to --Any--. In transparent firewall mode, you must set specific interfaces.

a.   From the Match Criteria: Original Packet > Source Interface drop-down list, choose the source interface.

b.   From the Match Criteria: Original Packet > Destination Interface drop-down list, choose the destination interface.

Step 3   Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the source interface network (the *real source address* and the *mapped destination address*). 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.

a. For the Match Criteria: Original Packet > Source Address, click the browse button and choose an existing network object or group or create a new object or group from the Browse Original Source Address dialog box. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. The default is **any**; only use this option when also setting the mapped address to **any**.

b. (Optional) For the Match Criteria: Original Packet > Destination Address, click the browse button ⋯ and choose an existing network object or group or create a new object or group from the Browse Original Destination Address dialog box.

Although the main feature of twice NAT is the inclusion of the destination IP address, the destination address is optional. If you do specify the destination address, you can configure static translation for that address or just use identity NAT for it. You might want to configure twice NAT without a destination address to take advantage of some of the other qualities of twice NAT, including the use of network object groups for real addresses, or manually ordering of rules. For more information, see Main Differences Between Network Object NAT and Twice NAT, page 5-15.

Step 4   (Optional) Identify the original packet source or destination port (the *real source port* **or** the *mapped destination port*). For the Match Criteria: Original Packet > Service, click the browse button and choose an existing TCP or UDP service object or create a new object from the Browse Original Service dialog box.

A service object can contain both a source and destination port. You should specify *either* the source *or* the destination port for both service objects. You should only specify *both* the source and destination ports if your application uses a fixed source port (such as some DNS servers); but fixed source ports are rare. In the rare case where you specify both the source and destination ports in the object, the original packet service object contains the real source port/mapped destination port; the translated packet service object contains the mapped source port/real 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. The "not equal" (**!=**) operator is not supported.

Step 5   Choose **Static** from the Match Criteria: Translated Packet > Source NAT Type drop-down list. Static is the default setting.

This setting only applies to the source address; the destination translation is always static.

Step 6   Identify the translated packet addresses; namely, the packet addresses as they appear on the destination interface network (the *mapped source address* and the *real destination address*). 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.
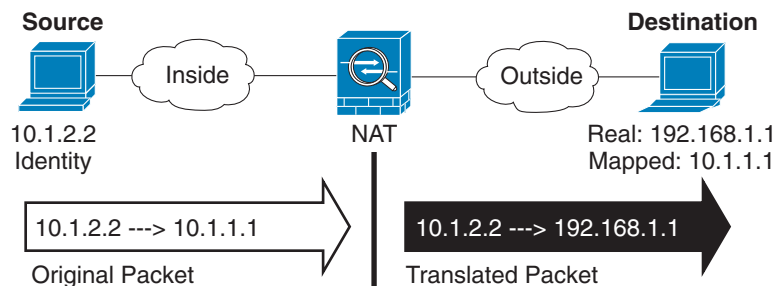
a.   For the Match Criteria: Translated Packet > Source Address, click the browse button and choose the same network object or group from the Browse Translated Source Address dialog box that you chose for the real source address. Use **any** if you specified **any** for the real address.

b.   For the Match Criteria: Translated Packet > Destination Address, click the browse button and choose an existing network object, group, or interface or create a new object or group from the Browse Translated Destination Address dialog box.

For identity NAT for the destination address, simply use the same object or group for both the real and mapped addresses.

If you want to translate the destination address, then the static mapping is typically one-to-one, so the real addresses have the same quantity as the mapped addresses. You can, however, have different quantities if desired. For more information, see Static NAT, page 5-3. See Guidelines and Limitations, page 7-2 for information about disallowed mapped IP addresses.

For static interface NAT with port translation only, choose an interface. If you specify an interface, be sure to also configure a a service translation. For more information, see Static Interface NAT with Port Translation, page 5-6.

Step 7   (Optional) Identify the translated packet source or destination port (the *mapped source port* **or** the *real destination port*). For the Match Criteria: Translated Packet > Service, click the browse button and choose an existing TCP or UDP service object or create a new object from the Browse Translated Service dialog box.

A service object can contain both a source and destination port. You should specify *either* the source *or* the destination port for both service objects. You should only specify *both* the source and destination ports if your application uses a fixed source port (such as some DNS servers); but fixed source ports are rare. In the rare case where you specify both the source and destination ports in the object, the original packet service object contains the real source port/mapped destination port; the translated packet service object contains the mapped source port/real 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. The "not equal" (**!=**) operator is not supported.

Step 8   (Optional) Configure NAT options in the Options area.

   a. Enable rule —Enables this NAT rule. The rule is enabled by default.

   b. Disable Proxy ARP on egress interface—Disables proxy ARP for incoming packets to the mapped IP addresses. See Mapped Addresses and Routing, page 5-22 for more information.

   c. (Routed mode; interface(s) specified) Lookup route table to locate egress interface—Determines the egress interface using a route lookup instead of using the interface specified in the NAT command. See Determining the Egress Interface, page 5-24 for more information.

   d. Direction—To make the rule unidirectional, choose **Unidirectional**. The default is Both. Making the rule unidirectional prevents traffic from initiating connections to the real addresses. You might want to use this setting for testing purposes.

   e. Description—Adds a description about the rule up to 200 characters in length.

Note Although the "Translate DNS replies that match this rule" check box is available if you do not configure a destination address, this option is not applicable to identity NAT because you are translating the address to itself, so the DNS reply does not need modification. See DNS and NAT, page 5-31 for more information.

Step 9 Click **OK**.

Configuring Per-Session PAT Rules

By default, all TCP PAT traffic and all UDP DNS traffic uses per-session PAT. To use multi-session PAT for traffic, you can configure per-session PAT rules: a permit rule uses per-session PAT, and a deny rule uses multi-session PAT. For more information about per-session vs. multi-session PAT, see Per-Session PAT vs. Multi-Session PAT (Version 9.0(1) and Later), page 5-11.

Detailed Steps

To configure a per-session PAT rule, see Configuring Per-Session PAT Rules, page 6-18.

Monitoring Twice NAT

The Monitoring > Properties > Connection Graphs > Xlates pane lets you view the active Network Address Translations in a graphical format. You can choose up to four types of statistics to show in one graph window. You can open multiple graph windows at the same time.

Fields

- Available Graphs—Lists the components you can graph.

    - Xlate Utilization—Displays the ASA NAT utilization.

- Graph Window Title—Shows the graph window name to which you want to add a graph type. To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title.

- Add—Click to move the selected entries in the Available Graphs list to the Selected Graphs list.

- Remove—Click to remove the selected entry from the Selected Graphs list.

- Show Graphs—Click to display a new or updated graph window.

The Monitoring > Properties > Connection Graphs > Perfmon pane lets you view the performance information in a graphical format. You can choose up to four types of statistics to show in one graph window. You can open multiple graph windows at the same time.

Fields

- Available Graphs—Lists the components you can graph.

    - AAA Perfmon—Displays the ASA AAA performance information.

    - Inspection Perfmon—Displays the ASA inspection performance information.

    - Web Perfmon—Displays the ASA web performance information, including URL access and URL server requests.

    - Connections Perfmon—Displays the ASA connections performance information.

    - Xlate Perfmon—Displays the ASA NAT performance information.

- Graph Window Title—Shows the graph window name to which you want to add a graph type. To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title.

- Add—Click to move the selected entries in the Available Graphs list to the Selected Graphs list.

- Remove—Click to remove the selected statistic type from the Selected Graphs list.

- Show Graphs—Click to display a new or updated graph window.

Configuration Examples for Twice NAT

This section includes the following configuration examples:

Different Translation Depending on the Destination (Dynamic PAT)

Figure 7-1 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*.

Figure 7-1   Twice NAT with Different Destination Addresses
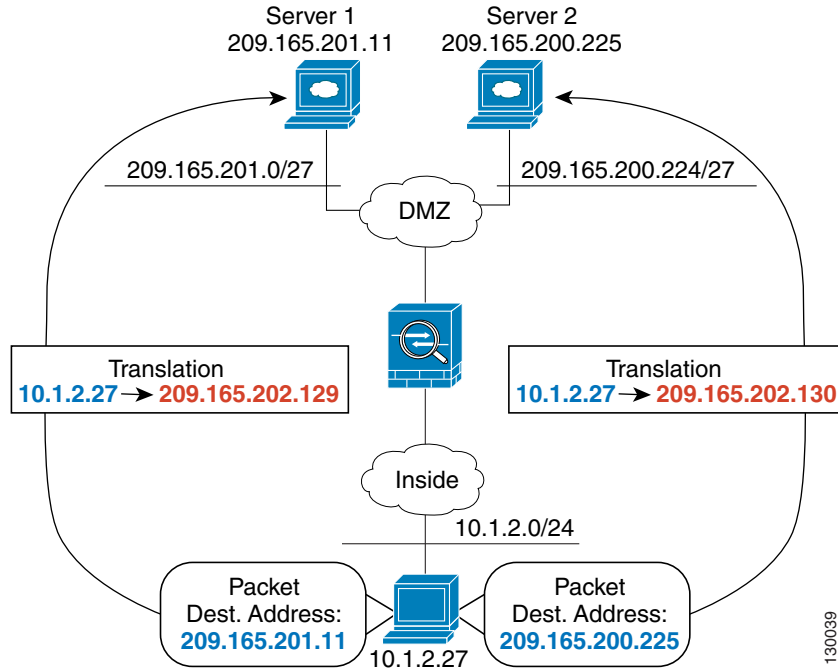
Step 1   Add a NAT rule for traffic from the inside network to DMZ network 1:

By default, the NAT rule is added to the end of section 1. If you want to add a NAT rule to section 3, after the network object NAT rules, choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.

Step 2   Set the source and destination interfaces:

Step 3   For the Original Source Address, click the browse button to add a new network object for the inside network in the Browse Original Source Address dialog box.

    a.   Add the new network object.

    b.   Define the inside network addresses, and click **OK**.

c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 4   For the Original Destination Address, click the browse button to add a new network object for DMZ network 1 in the Browse Original Destination Address dialog box.

a. Add the new network object.

b. Define the DMZ network 1 addresses, and click **OK**.

c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 5   Set the NAT Type to **Dynamic PAT (Hide)**:
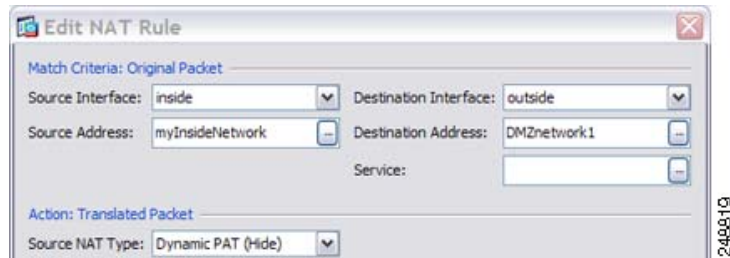
Step 6 For the Translated Source Address, click the browse button to add a new network object for the PAT address in the Browse Translated Source Address dialog box.

    a. Add the new network object.
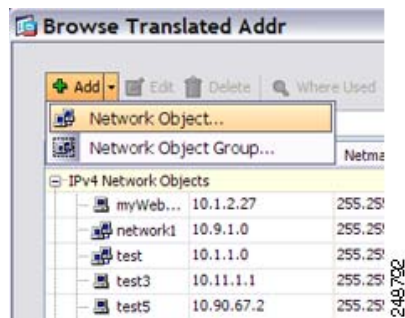
    b. Define the PAT address, and click **OK**.

    c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 7 For the Translated Destination Address, type the name of the Original Destination Address (DMZnetwork1) or click the browse button to choose it.

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.
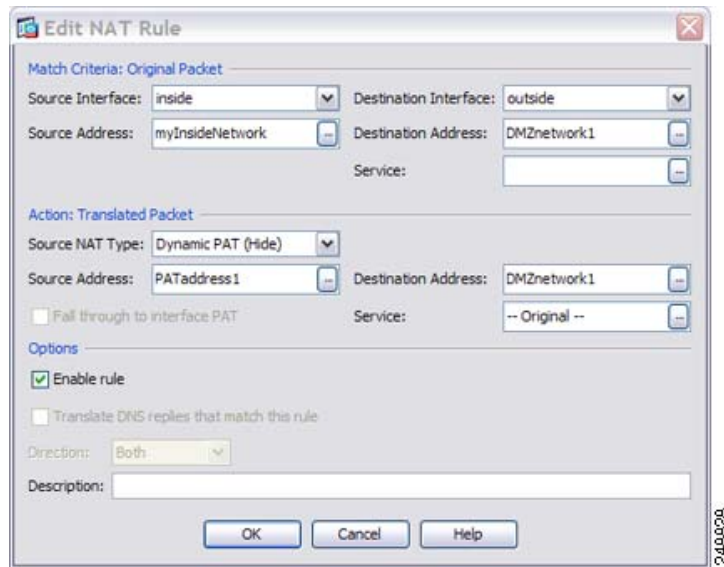
Step 8   Click **OK** to add the rule to the NAT table.

Step 9   Add a NAT rule for traffic from the inside network to DMZ network 2:
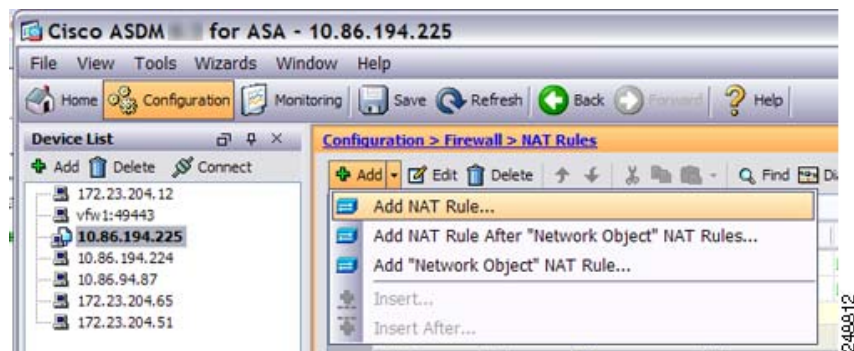
By default, the NAT rule is added to the end of section 1. If you want to add a NAT rule to section 3, after the network object NAT rules, choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.
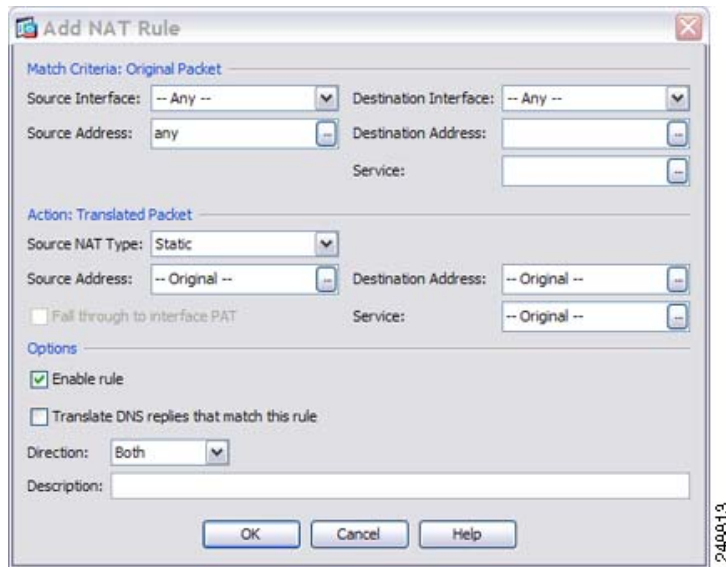
Step 10 Set the source and destination interfaces:

Step 11 For the Original Source Address, type the name of the inside network object (myInsideNetwork) or click the browse button to choose it.

Step 12 For the Original Destination Address, click the browse button to add a new network object for DMZ network 2 in the Browse Original Destination Address dialog box.

    a. Add the new network object.
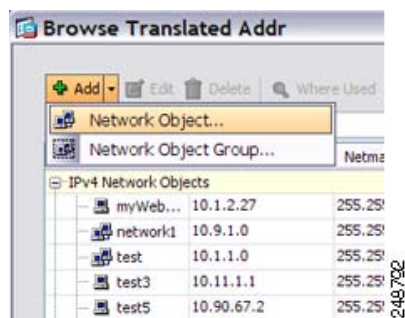
    b. Define the DMZ network 2 addresses, and click **OK**.

c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 13   Set the NAT Type to **Dynamic PAT (Hide)**:
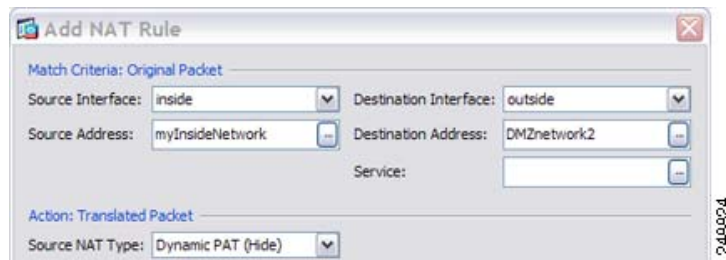
Step 14   For the Translated Source Address, click the browse button to add a new network object for the PAT address in the Browse Translated Source Address dialog box.

a.   Add the new network object.

b.   Define the PAT address, and click **OK**.

c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 15   For the Translated Destination Address, type the name of the Original Destination Address (DMZnetwork2) or click the browse button to choose it.

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.
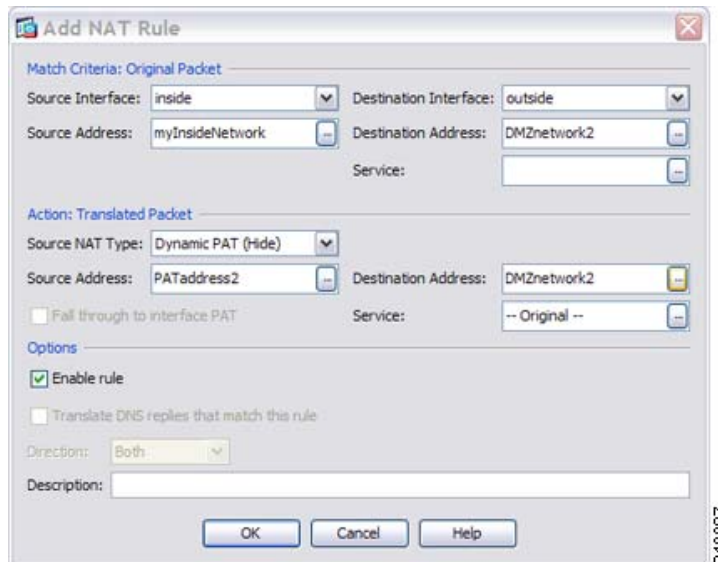
Step 16   Click **OK** to add the rule to the NAT table.

Step 17   Click **Apply**.

Different Translation Depending on the Destination Address and Port (Dynamic PAT)

Figure 7-2 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*.

Figure 7-2   *Twice NAT with Different Destination Ports*
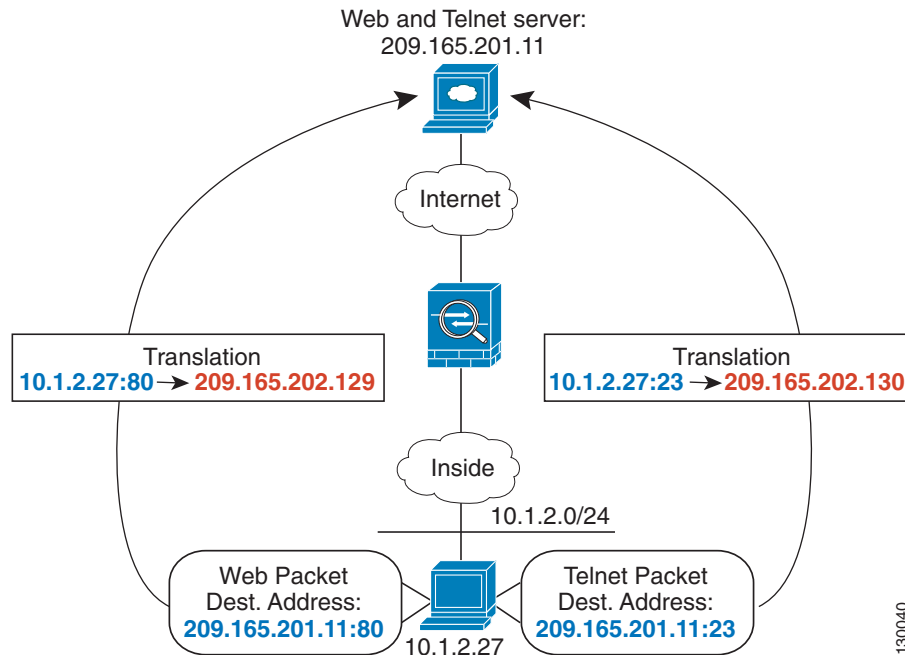
Step 1   Add a NAT rule for traffic from the inside network to the Telnet server:

By default, the NAT rule is added to the end of section 1. If you want to add a NAT rule to section 3, after the network object NAT rules, choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.

Step 2 Set the source and destination interfaces:

Step 3 For the Original Source Address, click the browse button to add a new network object for the inside network in the Browse Original Source Address dialog box.

    a. Add the new network object.

    b. Define the inside network addresses, and click **OK**.

c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 4 For the Original Destination Address, click the browse button to add a new network object for the Telnet/Web server in the Browse Original Destination Address dialog box.

   a. Add the new network object.
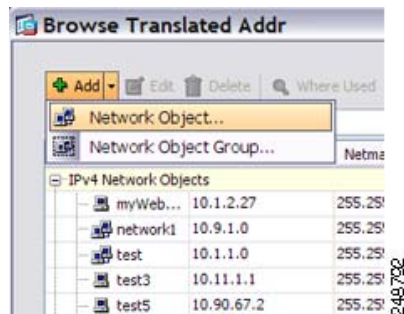
   b. Define the server address, and click **OK**.

   c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 5 For the Original Service, click the browse button to add a new service object for Telnet in the Browse Original Service dialog box.

  a. Add the new service object.

  b. Define the protocol and port, and click **OK**.

  c. Choose the new service object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 6 Set the NAT Type to **Dynamic PAT (Hide)**:
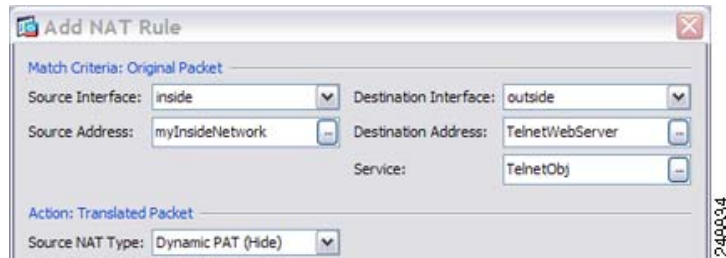
Step 7 For the Translated Source Address, click the browse button to add a new network object for the PAT address in the Browse Translated Source Address dialog box.

a. Add the new network object.
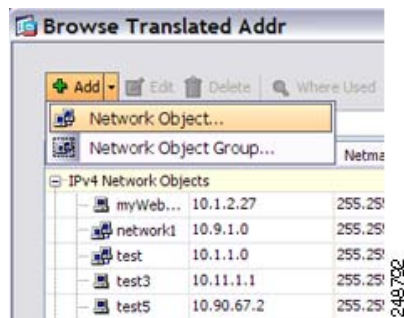
b. Define the PAT address, and click **OK**.

c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 8 For the Translated Destination Address, type the name of the Original Destination Address (TelnetWebServer) or click the browse button to choose it.

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.
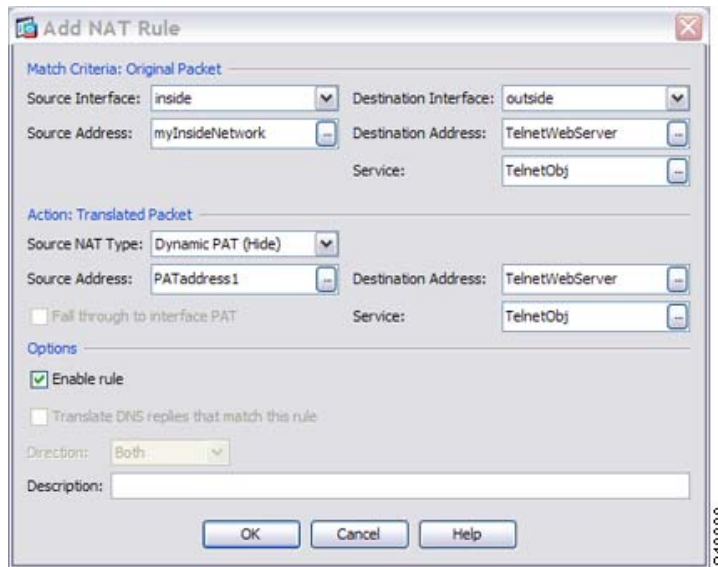
Step 9   Click **OK** to add the rule to the NAT table.

Step 10   Add a NAT rule for traffic from the inside network to the web server:
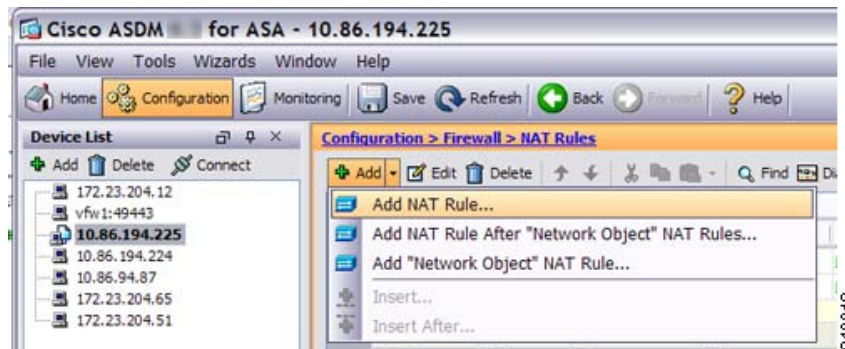
By default, the NAT rule is added to the end of section 1. If you want to add a NAT rule to section 3, after the network object NAT rules, choose **Add NAT Rule After Network Object NAT Rules**.

The Add NAT Rule dialog box appears.

Step 11 Set the real and mapped interfaces:

Step 12 For the Original Source Address, type the name of the inside network object (myInsideNetwork) or click the browse button to choose it.

Step 13 For the Original Destination Address, type the name of the Telnet/web server network object (TelnetWebServer) or click the browse button to choose it.

Step 14 For the Original Service, click the browse button to add a new service object for HTTP in the Browse Original Service dialog box.

   a. Add the new service object.

   b. Define the protocol and port, and click **OK**.

c. Choose the new service object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 15   Set the NAT Type to **Dynamic PAT (Hide)**:
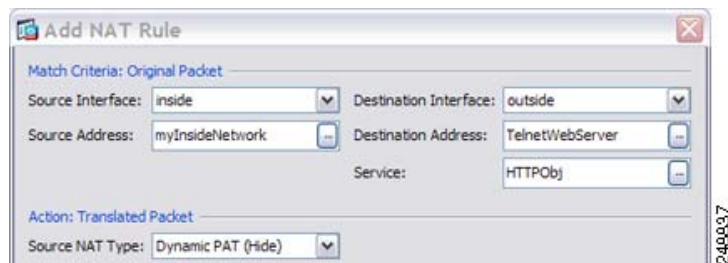
Step 16   For the Translated Source Address, click the browse button to add a new network object for the PAT address in the Browse Translated Source Address dialog box.

a. Add the new network object.
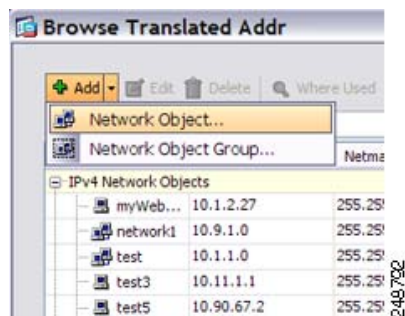
b. Define the PAT address, and click **OK**.

c. Choose the new network object by double-clicking it. Click **OK** to return to the NAT configuration.

Step 17 For the Translated Destination Address, type the name of the Original Destination Address (TelnetWebServer) or click the browse button to choose it.

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.
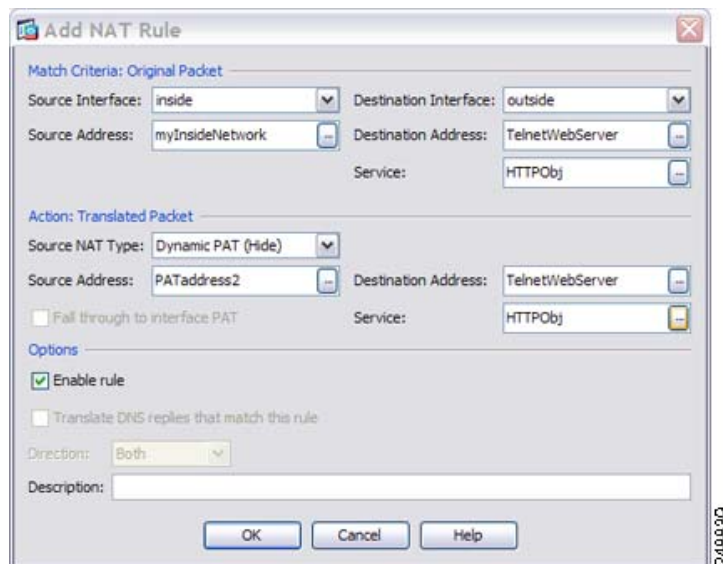
Step 18 Click **OK** to add the rule to the NAT table.

Step 19 Click **Apply**.

Feature History for Twice NAT

Table 7-1 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 7-1   Feature History for Twice NAT

Feature Name	Platform Releases	Feature Information
Twice NAT	8.3(1)	Twice NAT lets you identify both the source and destination address in a single rule. We modified the following screen: Configuration > Firewall > NAT Rules.
Identity NAT configurable proxy ARP and route lookup	8.4(2)/8.5(1)	In earlier releases for identity NAT, proxy ARP was disabled, and a route lookup was always used to determine the egress interface. You could not configure these settings. In 8.4(2) and later, the default behavior for identity NAT was changed to match the behavior of other static NAT configurations: proxy ARP is enabled, and the NAT configuration determines the egress interface (if specified) by default. You can leave these settings as is, or you can enable or disable them discretely. Note that you can now also disable proxy ARP for regular static NAT. For pre-8.3 configurations, the migration of NAT exempt rules (the **nat 0 access-list** command) to 8.4(2) and later now includes the following keywords to disable proxy ARP and to use a route lookup: **no-proxy-arp** and **route-lookup**. The **unidirectional** keyword that was used for migrating to 8.3(2) and 8.4(1) is no longer used for migration. When upgrading to 8.4(2) from 8.3(1), 8.3(2), and 8.4(1), all identity NAT configurations will now include the **no-proxy-arp** and **route-lookup** keywords, to maintain existing functionality. The **unidirectional** keyword is removed. We modified the following screen: Configuration > Firewall > NAT Rules > Add/Edit NAT Rule
PAT pool and round robin address assignment	8.4(2)/8.5(1)	You can now specify a pool of PAT addresses instead of a single address. You can also optionally enable round-robin assignment of PAT addresses instead of first using all ports on a PAT address before using the next address in the pool. These features help prevent a large number of connections from a single PAT address from appearing to be part of a DoS attack and makes configuration of large numbers of PAT addresses easy. We modified the following screens: Configuration > Firewall > NAT Rules > Add/Edit NAT Rule.

Table 7-1 ***Feature History for Twice NAT (continued)***

Feature Name	Platform Releases	Feature Information
Round robin PAT pool allocation uses the same IP address for existing hosts	8.4(3)	When using a PAT pool with round robin allocation, if a host has an existing connection, then subsequent connections from that host will use the same PAT IP address if ports are available. We did not modify any screens. *This feature is not available in 8.5(1) or 8.6(1).*
Flat range of PAT ports for a PAT pool	8.4(3)	If available, the real source port number is used for the mapped port. However, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 0 to 511, 512 to 1023, and 1024 to 65535. Therefore, ports below 1024 have only a small PAT pool. If you have a lot of traffic that uses the lower port ranges, when using a PAT pool, you can now specify a flat range of ports to be used instead of the three unequal-sized tiers: either 1024 to 65535, or 1 to 65535. We modified the following screens: Configuration > Firewall > NAT Rules > Add/Edit NAT Rule. *This feature is not available in 8.5(1) or 8.6(1).*
Extended PAT for a PAT pool	8.4(3)	Each PAT IP address allows up to 65535 ports. If 65535 ports do not provide enough translations, you can now enable extended PAT for a PAT pool. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. We modified the following screens: Configuration > Firewall > NAT Rules > Add/Edit NAT Rule. *This feature is not available in 8.5(1) or 8.6(1).*

Table 7-1 *Feature History for Twice NAT (continued)*

Feature Name	Platform Releases	Feature Information
Automatic NAT rules to translate a VPN peer's local IP address back to the peer's real IP address	8.4(3)	In rare situations, you might want to use a VPN peer's real IP address on the inside network instead of an assigned local IP address. Normally with VPN, the peer is given an assigned local IP address to access the inside network. However, you might want to translate the local IP address back to the peer's real public IP address if, for example, your inside servers and network security is based on the peer's real IP address.
		You can enable this feature on one interface per tunnel group. Object NAT rules are dynamically added and deleted when the VPN session is established or disconnected. You can view the rules using the **show nat** command.
		Note Because of routing issues, we do not recommend using this feature unless you know you need this feature; contact Cisco TAC to confirm feature compatibility with your network. See the following limitations:
		• Only supports Cisco IPsec and AnyConnect Client.
		• Return traffic to the public IP addresses must be routed back to the ASA so the NAT policy and VPN policy can be applied.
		• Does not support load-balancing (because of routing issues).
		• Does not support roaming (public IP changing).
		ASDM does not support this command; enter the command using the Command Line Tool.
NAT support for IPv6	9.0(1)	NAT now supports IPv6 traffic, as well as translating between IPv4 and IPv6. Translating between IPv4 and IPv6 is not supported in transparent mode.
		We modified the following screen: Configuration > Firewall > NAT Rules.

Table 7-1   *Feature History for Twice NAT (continued)*

Feature Name	Platform Releases	Feature Information
NAT support for reverse DNS lookups	9.0(1)	NAT now supports translation of the DNS PTR record for reverse DNS lookups when using IPv4 NAT, IPv6 NAT, and NAT64 with DNS inspection enabled for the NAT rule.
Per-session PAT	9.0(1)	The per-session PAT feature improves the scalability of PAT and, for clustering, allows each member unit to own PAT connections; multi-session PAT connections have to be forwarded to and owned by the master unit. At the end of a per-session PAT session, the ASA sends a reset and immediately removes the xlate. This reset causes the end node to immediately release the connection, avoiding the TIME_WAIT state. Multi-session PAT, on the other hand, uses the PAT timeout, by default 30 seconds. For "hit-and-run" traffic, such as HTTP or HTTPS, the per-session feature can dramatically increase the connection rate supported by one address. Without the per-session feature, the maximum connection rate for one address for an IP protocol is approximately 2000 per second. With the per-session feature, the connection rate for one address for an IP protocol is 65535/*average-lifetime*. By default, all TCP traffic and UDP DNS traffic use a per-session PAT xlate. For traffic that requires multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT by creating a per-session deny rule. We introduced the following screen: Configuration > Firewall > Advanced > Per-Session NAT Rules.

P ART 3

Application Inspection

Getting Started with Application Layer Protocol Inspection

This chapter describes how to configure application layer protocol inspection. Inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the ASA to do a deep packet inspection instead of passing the packet through the fast path (see the general operations configuration guide for more information about the fast path). As a result, inspection engines can affect overall throughput. Several common inspection engines are enabled on the ASA by default, but you might need to enable others depending on your network.

This chapter includes the following sections:

Information about Application Layer Protocol Inspection

This section includes the following topics:

How Inspection Engines Work

As illustrated in Figure 8-1, the ASA uses three databases for its basic operation:

- ACLs—Used for authentication and authorization of connections based on specific networks, hosts, and services (TCP/UDP port numbers).
- Inspections—Contains a static, predefined set of application-level inspection functions.
- Connections (XLATE and CONN tables)—Maintains state and other information about each established connection. This information is used by the Adaptive Security Algorithm and cut-through proxy to efficiently forward traffic within established sessions.

Figure 8-1   How Inspection Engines Work

In Figure 8-1, operations are numbered in the order they occur, and are described as follows:

1. A TCP SYN packet arrives at the ASA to establish a new connection.

2. The ASA checks the ACL database to determine if the connection is permitted.

3. The ASA creates a new entry in the connection database (XLATE and CONN tables).

4. The ASA checks the Inspections database to determine if the connection requires application-level inspection.

5. After the application inspection engine completes any required operations for the packet, the ASA forwards the packet to the destination system.

6. The destination system responds to the initial request.

7. The ASA receives the reply packet, looks up the connection in the connection database, and forwards the packet because it belongs to an established session.

The default configuration of the ASA includes a set of application inspection entries that associate supported protocols with specific TCP or UDP port numbers and that identify any special handling required.

When to Use Application Protocol Inspection

When a user establishes a connection, the ASA checks the packet against ACLs, creates an address translation, and creates an entry for the session in the fast path, so that further packets can bypass time-consuming checks. However, the fast path relies on predictable port numbers and does not perform address translations inside a packet.

Many protocols open secondary TCP or UDP ports. The initial session on a well-known port is used to negotiate dynamically assigned port numbers.

Other applications embed an IP address in the packet that needs to match the source address that is normally translated when it goes through the ASA.

If you use applications like these, then you need to enable application inspection.

When you enable application inspection for a service that embeds IP addresses, the ASA translates embedded addresses and updates any checksum or other fields that are affected by the translation.

When you enable application inspection for a service that uses dynamically assigned ports, the ASA monitors sessions to identify the dynamic port assignments, and permits data exchange on these ports for the duration of the specific session.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Failover Guidelines

State information for multimedia sessions that require inspection are not passed over the state link for stateful failover. The exception is GTP, which is replicated over the state link.

IPv6 Guidelines

Supports IPv6 for the following inspections:

- DNS
- FTP
- HTTP
- ICMP
- SIP
- SMTP
- IPsec pass-through
- IPv6

Supports NAT64 for the following inspections:

- DNS
- FTP
- HTTP
- ICMP

Additional Guidelines and Limitations

Some inspection engines do not support PAT, NAT, outside NAT, or NAT between same security interfaces. See Default Settings and NAT Limitations, page 8-4 for more information about NAT support.

For all the application inspections, the ASA limits the number of simultaneous, active data connections to 200 connections. For example, if an FTP client opens multiple secondary connections, the FTP inspection engine allows only 200 active connections and the 201 connection is dropped and the adaptive security appliance generates a system error message.

Inspected protocols are subject to advanced TCP-state tracking, and the TCP state of these connections is not automatically replicated. While these connections are replicated to the standby unit, there is a best-effort attempt to re-establish a TCP state.

Default Settings and NAT Limitations

By default, the configuration includes a policy that matches all default application inspection traffic and applies inspection to the traffic on all interfaces (a global policy). Default application inspection traffic includes traffic to the default ports for each protocol. You can only apply one global policy, so if you want to alter the global policy, for example, to apply inspection to non-standard ports, or to add inspections that are not enabled by default, you need to either edit the default policy or disable it and apply a new one.

Table 8-1 lists all inspections supported, the default ports used in the default class map, and the inspection engines that are on by default, shown in bold. This table also notes any NAT limitations.

Table 8-1   Supported Application Inspection Engines

Application[1]	Default Port	NAT Limitations	Standards[2]	Comments
CTIQBE	TCP/2748	No extended PAT. No NAT64. (Clustering) No static PAT.	—	—
DCERPC	TCP/135	No NAT64.	—	—
DNS over UDP	UDP/53	No NAT support is available for name resolution through WINS.	RFC 1123	—
FTP	TCP/21	(Clustering) No static PAT.	RFC 959	—
GTP	UDP/3386 UDP/2123	No extended PAT. No NAT64.	—	Requires a special license.
H.323 H.225 and **RAS**	TCP/1720 UDP/1718 UDP (RAS) 1718-1719	No dynamic NAT or PAT. Static PAT may not work. (Clustering) No static PAT. No extended PAT. No per-session PAT. No NAT on same security interfaces. No outside NAT. No NAT64.	ITU-T H.323, H.245, H225.0, Q.931, Q.932	—
HTTP	TCP/80	—	RFC 2616	Beware of MTU limitations stripping ActiveX and Java. If the MTU is too small to allow the Java or ActiveX tag to be included in one packet, stripping may not occur.
ICMP	—	—	—	—

Table 8-1   Supported Application Inspection Engines (continued)

Application[1]	Default Port	NAT Limitations	Standards[2]	Comments
ICMP ERROR	—	—	—	—
ILS (LDAP)	TCP/389	No extended PAT. No NAT64.	—	—
Instant Messaging (IM)	Varies by client	No extended PAT. No NAT64.	RFC 3860	—
IP Options	—	No NAT64.	RFC 791, RFC 2113	—
IPsec Pass Through	UDP/500	No PAT. No NAT64.	—	—
IPv6	—	No NAT64.	RFC 2460	—
MGCP	UDP/2427, 2727	No extended PAT. No NAT64. (Clustering) No static PAT.	RFC 2705bis-05	—
MMP	TCP 5443	No extended PAT. No NAT64.	—	—
NetBIOS Name Server over IP	UDP/137, 138 (Source ports)	No extended PAT. No NAT64.	—	NetBIOS is supported by performing NAT of the packets for NBNS UDP port 137 and NBDS UDP port 138.
PPTP	TCP/1723	No NAT64. (Clustering) No static PAT.	RFC 2637	—
RADIUS Accounting	1646	No NAT64.	RFC 2865	—
RSH	TCP/514	No PAT. No NAT64. (Clustering) No static PAT.	Berkeley UNIX	—
RTSP	TCP/554	No extended PAT. No outside NAT. No NAT64. (Clustering) No static PAT.	RFC 2326, 2327, 1889	No handling for HTTP cloaking.
ScanSafe (Cloud Web Security)	TCP/80 TCP/413	—	—	These ports are not included in the default-inspection-traffic class for the ScanSafe inspection.

Table 8-1   Supported Application Inspection Engines (continued)

Application[1]	Default Port	NAT Limitations	Standards[2]	Comments
SIP	TCP/5060 UDP/5060	No outside NAT. No NAT on same security interfaces. No extended PAT. No per-session PAT. No NAT64. (Clustering) No static PAT.	RFC 2543	—
SKINNY (SCCP)	TCP/2000	No outside NAT. No NAT on same security interfaces. No extended PAT. No per-session PAT. No NAT64. (Clustering) No static PAT.	—	Does not handle TFTP uploaded Cisco IP Phone configurations under certain circumstances.
SMTP and **ESMTP**	TCP/25	No NAT64.	RFC 821, 1123	—
SNMP	UDP/161, 162	No NAT or PAT.	RFC 1155, 1157, 1212, 1213, 1215	v.2 RFC 1902-1908; v.3 RFC 2570-2580.
SQL*Net	TCP/1521	No extended PAT. No NAT64. (Clustering) No static PAT.	—	v.1 and v.2.
Sun RPC over UDP and TCP	UDP/111	No extended PAT. No NAT64.	—	The default rule includes UDP port 111; if you want to enable Sun RPC inspection for TCP port 111, you need to create a new rule that matches TCP port 111 and performs Sun RPC inspection.
TFTP	UDP/69	No NAT64. (Clustering) No static PAT.	RFC 1350	Payload IP addresses are not translated.
WAAS	—	No extended PAT. No NAT64.	—	—
XDCMP	UDP/177	No extended PAT. No NAT64. (Clustering) No static PAT.	—	—

1. Inspection engines that are enabled by default for the default port are in bold.

2. The ASA is in compliance with these standards, but it does not enforce compliance on packets being inspected. For example, FTP commands are supposed to be in a particular order, but the ASA does not enforce the order.

Configuring Application Layer Protocol Inspection

This feature uses Security Policy Rules to create a service policy. Service policies provide a consistent and flexible way to configure ASA features. For example, you can use a service policy to create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications. See Chapter 1, "Service Policy," for more information.

Inspection is enabled by default for some applications. See Default Settings and NAT Limitations, page 8-4 section for more information. Use this section to modify your inspection policy.

Detailed Steps

*

Step 1   Choose **Configuration > Firewall > Service Policy Rules**.

Step 2   Add or edit a service policy rule according to the Adding a Service Policy Rule for Through Traffic, page 1-8.

If you want to match non-standard ports, then create a new rule for the non-standard ports. See Default Settings and NAT Limitations, page 8-4 for the standard ports for each inspection engine. You can combine multiple rules in the same service policy if desired, so you can create one rule to match certain traffic, and another to match different traffic. However, if traffic matches a rule that contains an inspection action, and then matches another rule that also has an inspection action, only the first matching rule is used.

Step 3   In the Edit Service Policy Rule > Rule Actions dialog box, click the **Protocol Inspection** tab.

For a new rule, the dialog box is called Add Service Policy Rule Wizard - Rule Actions.

Step 4   Select each inspection type that you want to apply.

Step 5   (Optional) Some inspection engines let you control additional parameters when you apply the inspection to the traffic. Click **Configure** for each inspection type to configure an inspect map.

You can either choose an existing map, or create a new one. You can predefine inspect maps in the Configuration > Firewall > Objects > Inspect Maps pane.

Step 6   You can configure other features for this rule if desired using the other Rule Actions tabs.

Step 7   Click **OK** (or **Finish** from the wizard).

Inspection of Basic Internet Protocols

This chapter describes how to configure application layer protocol inspection. Inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the ASA to do a deep packet inspection instead of passing the packet through the fast path. As a result, inspection engines can affect overall throughput.

Several common inspection engines are enabled on the ASA by default, but you might need to enable others depending on your network.

This chapter includes the following sections:

DNS Inspection

This section describes DNS application inspection. This section includes the following topics:

Information About DNS Inspection

General Information About DNS

A single connection is created for multiple DNS sessions, as long as they are between the same two hosts, and the sessions have the same 5-tuple (source/destination IP address, source/destination port, and protocol). DNS identification is tracked by app_id, and the idle timer for each app_id runs independently. Because the app_id expires independently, a legitimate DNS response can only pass through the ASA within a limited period of time and there is no resource build-up.

DNS Inspection Actions

DNS inspection is enabled by default. You can customize DNS inspection to perform many tasks:

- Translate the DNS record based on the NAT configuration. For more information, see DNS and NAT, page 5-31.
- Enforce message length, domain-name length, and label length.
- Verify the integrity of the domain-name referred to by the pointer if compression pointers are encountered in the DNS message.
- Check to see if a compression pointer loop exists.
- Inspect packets based on the DNS header, type, class and more.

Default Settings for DNS Inspection

DNS inspection is enabled by default, using the preset_dns_map inspection class map:

- The maximum DNS message length is 512 bytes.
- The maximum client DNS message length is automatically set to match the Resource Record.
- DNS Guard is enabled, so the ASA tears down the DNS session associated with a DNS query as soon as the DNS reply is forwarded by the ASA. The ASA also monitors the message exchange to ensure that the ID of the DNS reply matches the ID of the DNS query.
- Translation of the DNS record based on the NAT configuration is enabled.
- Protocol enforcement is enabled, which enables DNS message format check, including domain name length of no more than 255 characters, label length of 63 characters, compression, and looped pointer check.

(Optional) Configuring a DNS Inspection Policy Map and Class Map

To match DNS packets with certain characteristics and perform special actions, create a DNS inspection policy map. You can also configure a DNS inspection class map to group multiple match criteria for reference within the inspection policy map. You can then apply the inspection policy map when you enable DNS inspection.

Prerequisites

If you want to match a DNS message domain name list, then create a regular expression using one of the methods below:

- Creating a regular expression (see the general operations configuration guide).
- Creating a regular expression class map (see the general operations configuration guide).

Detailed Steps

Step 1   Choose **Configuration > Firewall > Objects > Inspect Maps > DNS**.

The Configure DNS Maps pane appears.

Step 2   Click **Add**.

The Add IPv6 Inspection Map dialog box appears.

Step 3   In the Name field, name the inspection policy map.

Step 4   (Optional) In the Description field, add a description.

Step 5   Do one of the following:

- To use one of the preset security levels (Low, Medium, or High), drag the Security Level knob, then click **OK** to add the inspection policy map. You can skip the rest of this procedure.
- To customize each parameter and/or to configure packet matching inspection, click **Details**.

Detailed Steps—Protocol Conformance

Step 1   Configure the following Protocol Conformance parameters:
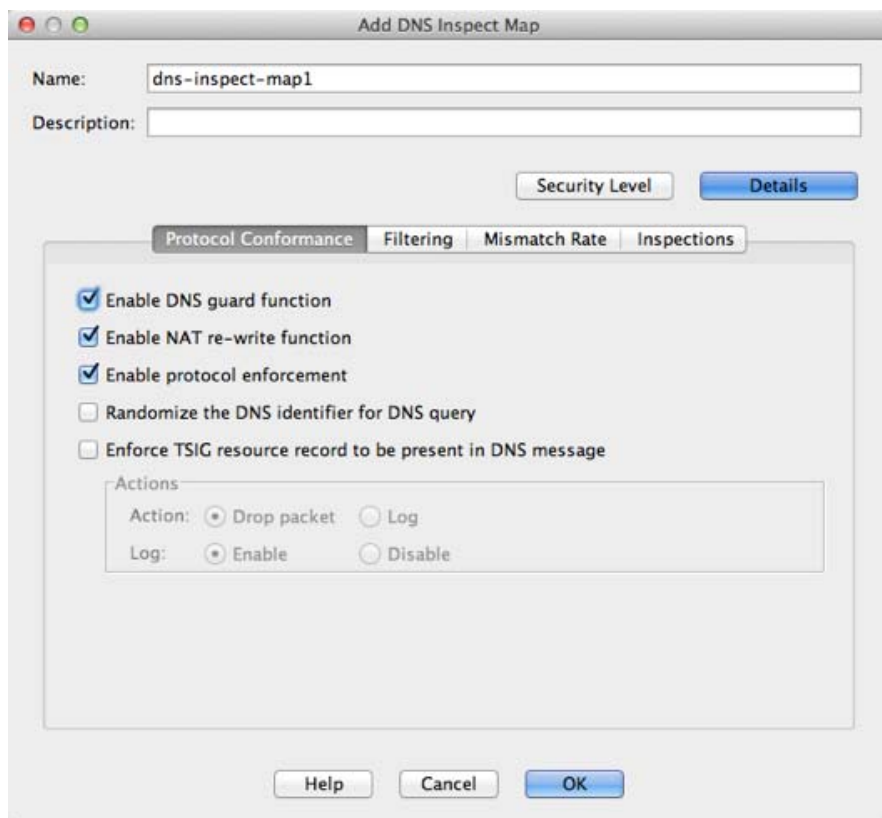
Step 2   **Enable DNS guard function**—Enables DNS Guard. The ASA tears down the DNS session associated with a DNS query as soon as the DNS reply is forwarded by the ASA. The ASA also monitors the message exchange to ensure that the ID of the DNS reply matches the ID of the DNS query.

Step 3   **Enable NAT re-write function**—Translates the DNS record based on the NAT configuration.

Step 4   **Enable protocol enforcement**—Enables DNS message format check, including domain name length of no more than 255 characters, label length of 63 characters, compression, and looped pointer check.

Step 5   **Randomize the DNS identifier for DNS query**—Randomizes the DNS identifier for a DNS query.

Step 6   **Enforce TSIG resource record to be present in DNS message**—Requires a TSIG resource record to be present. Actions include:

   • Action: **Drop packet** or **Log**—Drop or log a non-conforming packet.

   • Log: **Enable** or **Disable**—If you selected Drop packet, you can also enable logging.

Detailed Steps—Filtering

Step 1   Click the **Filtering** tab.
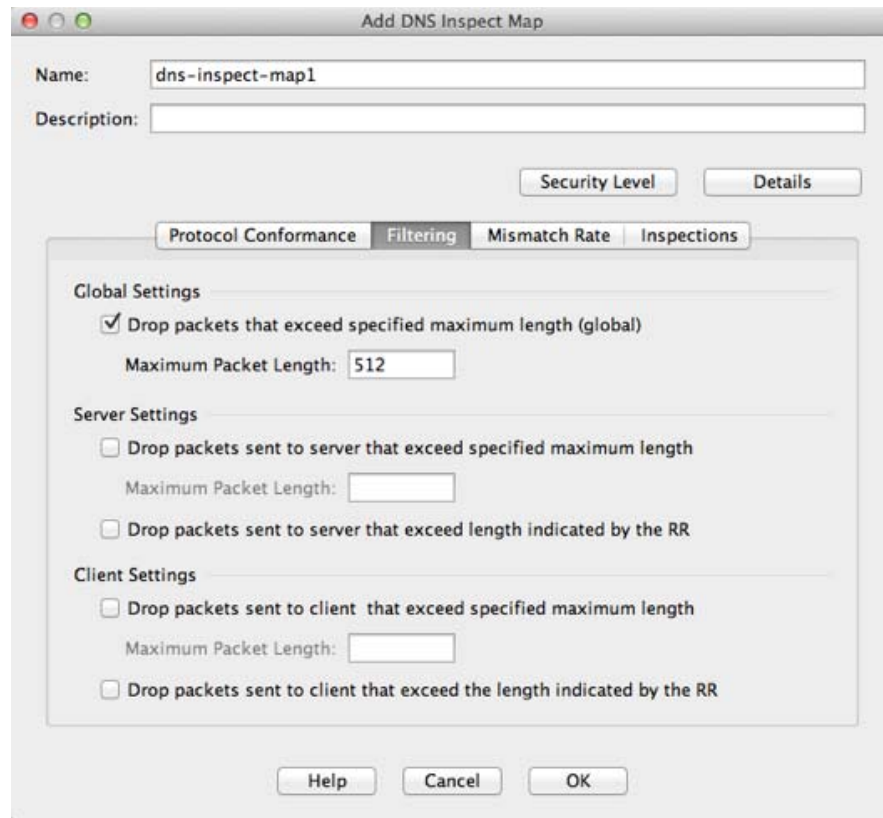
Step 2 Global Settings: **Drop packets that exceed specified maximum length (global)**—Sets the maximum DNS message length, from 512 to 65535 bytes.

Step 3 Server Settings: **Drop packets that exceed specified maximum length** and **Drop packets sent to server that exceed length indicated by the RR**—Sets the maximum server DNS message length, from 512 to 65535 bytes, or sets the maximum length to the value in the Resource Record. If you enable both settings, the lower value is used.

Step 4 Client Settings: **Drop packets that exceed specified maximum length** and **Drop packets sent to server that exceed length indicated by the RR**—Sets the maximum client DNS message length, from 512 to 65535 bytes, or sets the maximum length to the value in the Resource Record. If you enable both settings, the lower value is used.

Detailed Steps—Mismatch Rate

Step 1 Click the **Mismatch Rate** tab.

Step 2   **Enable logging when DNS ID mismatch rate exceeds specified rate**—Enables logging for excessive
DNS ID mismatches, where the Mismatch Instance Threshold and Time Interval fields specify the
maximum number of mismatch instances per *x* seconds before a system message log is sent.

Detailed Steps—Inspections

Step 1   Click the **Inspections** tab.

Step 2   Click **Add**.

The Add DNS Inspect dialog box appears.
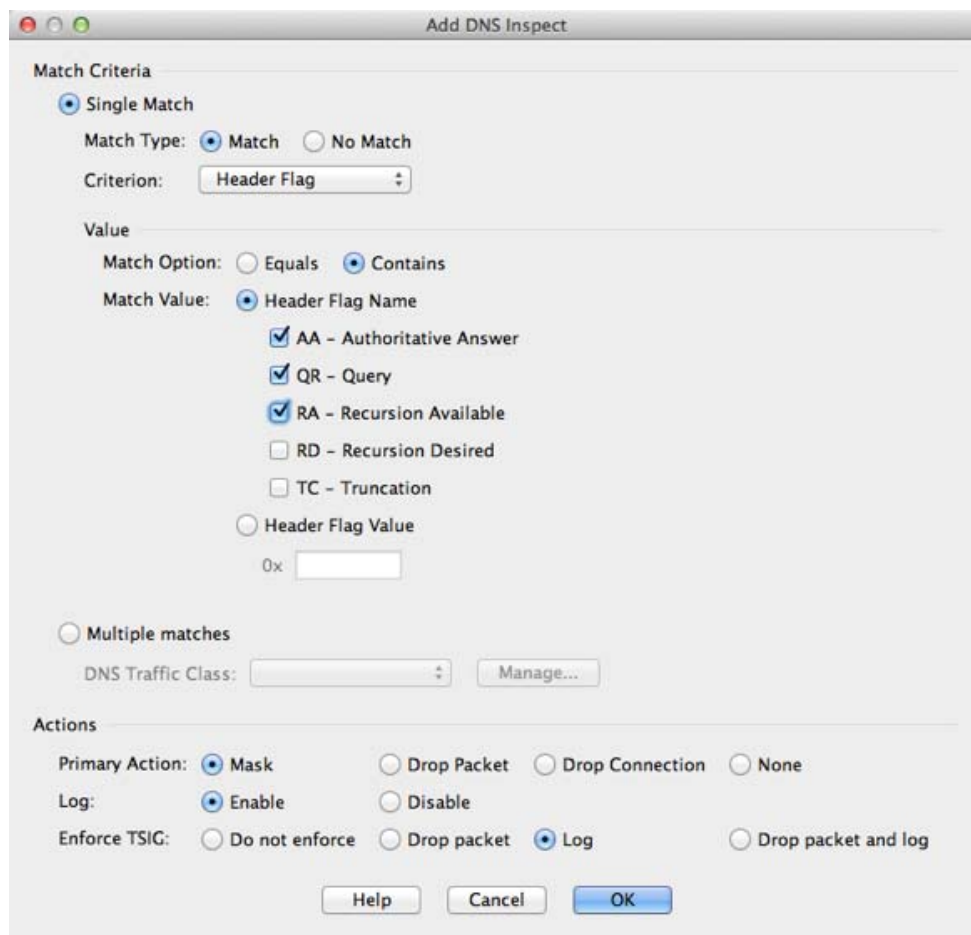
Step 3   You can configure DNS inspections using the following methods:

- **Single Match**—Match a single criterion, and identify the action for the match.

- **Multiple matches**—Match multiple criteria by creating an inspection class map.

The difference between creating a class map and defining the traffic match directly in the inspection policy map is that the class map lets you create more complex match criteria, and you can reuse class maps. If you want different actions for each criteria, use the single match option; you can only set one action for the entire class map.

You can add multiple class maps and single matches in the same policy map.

Actions for each Single Match, or for a Multiple match class map include:

- Primary Action:
    - Mask
    - Drop Packet
    - Drop Connection
    - None

- Log:
    - Enable
    - Disable

- Enforce TSIG: Requires a TSIG resource record to be present.
  - Do not enforce
  - Drop packet
  - Log
  - Drop packet and log

Not all combinations are valid for all matching criteria. For example, you can configure both Mask and Enforce TSIG together only for the Criterion: Header Flag option.

Step 4   For Multiple matches, if you predefined a class map on the Configuration > Firewall > Objects > Class Maps > DNS pane, you can select it from the drop-down list, set the Actions, and click **OK**.

To add a new class map:

a.   Click **Manage**.

The Manage DNS Class Maps dialog box appears

b.   Click **Add**.

The Add DNS Traffic Class Map dialog box appears.

c.   Click **Add**.

The Add DNS Match Criterion dialog box appears.

The match criteria are the same for a class map or for single matches; the following steps apply to both methods. The only difference is that you do not set an Action for each criterion in a class map.

Step 5   From the Criterion drop-down list, choose one of the following criteria:

- **Header Flag**:

Set the following Value parameters:

- Match Option: **Equals** or **Contains**. If you choose Header Flag Name, and check multiple flags, you can set the ASA to match a packet only if all flags are present (Equals) or if any one of the flags is present (Contains).

- Match Value: **Header Flag Name** or **Header Flag Value**. If you click **Header Flag Name**, you can check one or more well-known flag values. If you want to specify a hex value, click the **Header Flag Value** radio button, and enter the hex value in the field.

- **Type**:
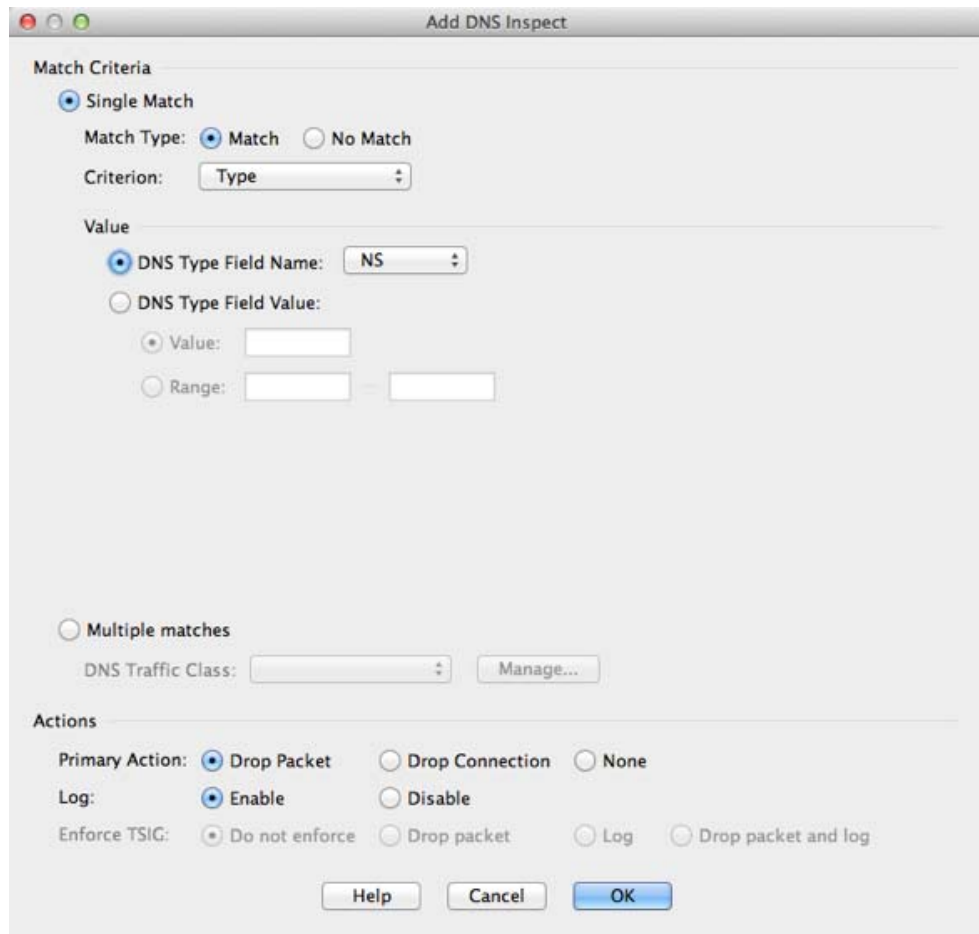
Set the following Value parameters:

- **DNS Type Field Name**—Lists the DNS types to select.

    A—IPv4 address

    AXFR—Full (zone) transfer

    CNAME—Canonical name

    IXFR—Incremental (zone) transfer

    NS—Authoritative name server

    SOA—Start of a zone of authority

    TSIG—Transaction signature

- **DNS Type Field Value:**

    Value—Lets you enter a value between 0 and 65535 to match.

    Range—Lets you enter a range match. Both values between 0 and 65535.

- **Class**:

Set the following Value parameters:

- **DNS Class Field Name: Internet**—Internet is the only option.
- **DNS Class Field Value**:

    Value—Lets you enter a value between 0 and 65535.

    Range—Lets you enter a range match. Both values between 0 and 65535.

- **Question**: Matches a DNS question.

- **Resource Record**:
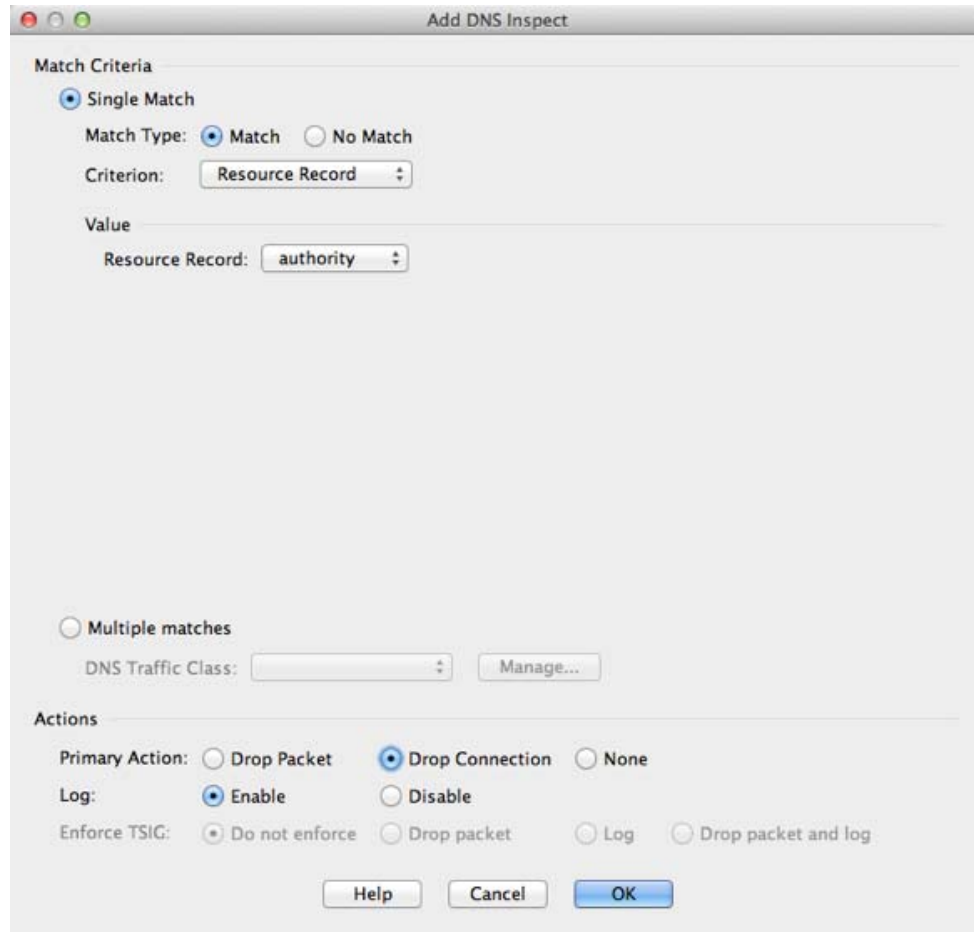
Set the following Value parameters:

– Resource Record:

additional—DNS additional resource record

answer—DNS answer resource record

authority—DNS authority resource record

• **Domain Name**:
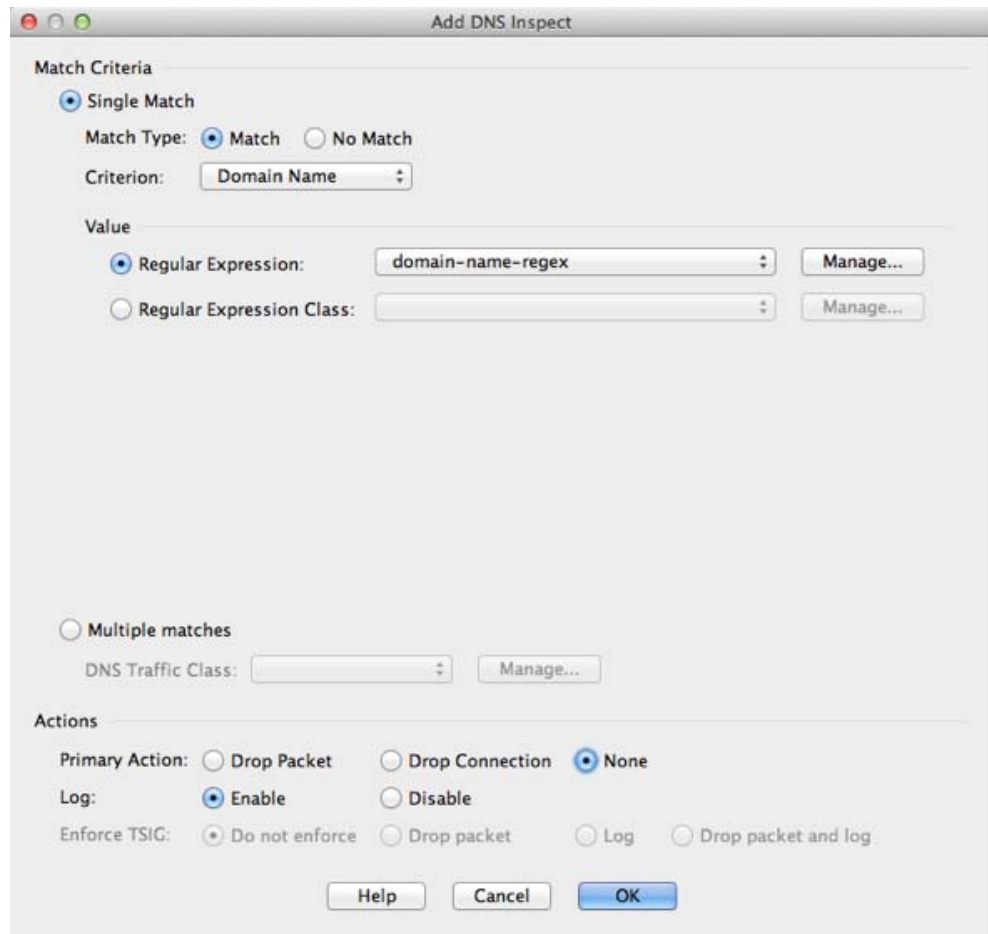
Set the following Value parameters:

- **Regular Expression**—Choose an existing regular expression from the drop-down menu, or click **Manage** to add a new one. See the general operations configuration guide.

- **Regular Expression Class**—Choose an existing regular expression class map from the drop-down menu, or click **Manage** to add a new one. See the general operations configuration guide.

Step 6 For a class map:

  a. Click **OK** to add the match to the map.

  b. Add more matches as desired.

  c. Click **OK** to finish the class map.

  d. Click **OK** to return to the Add DNS Inspect Map dialog box.

Step 7 Set the action for the Single Match, or for the Multiple matches class map; see Step 3 for actions.

Step 8 Click **OK** to return to the Add DNS Inspect dialog box.

Step 9 In some cases when you have more than one match in the inspection policy map, you can order the matches using the Move Up and Move Down buttons. Generally, the order is determined by internal ASA rules, so these buttons are not available for most entries. However, if you have a direct match and a class map that have the same match, then the order in the configuration determines which match is used, so these buttons are enabled. See Guidelines and Limitations, page 2-2 for more information.

Step 10   Click **OK** to save the DNS inspect map.

Step 11   Click **Apply**.

Configuring DNS Inspection

The default ASA configuration includes many default inspections on default ports applied globally on all interfaces. A common method for customizing the inspection configuration is to customize the default global policy. The steps in this section show how to edit the default global policy, but you can alternatively create a new service policy as desired, for example, an interface-specific policy.

Detailed Steps

Step 1   Configure a service policy on the Configuration > Firewall > Service Policy Rules pane according to Chapter 1, "Service Policy."

You can configure DNS inspection as part of a new service policy rule, or you can edit an existing service policy.

Step 2   On the Rule Actions dialog box, click the **Protocol Inspections** tab.

Step 3   (To change an in-use policy) If you are editing any in-use policy to use a different DNS inspection policy map, you must disable the DNS inspection, and then re-enable it with the new DNS inspection policy map name:

   a.   Uncheck the **DNS** check box.

   b.   Click **OK**.

   c.   Click **Apply**.

   d.   Repeat these steps to return to the Protocol Inspections tab.

Step 4   Check the **DNS** check box.

Step 5   Click **Configure**.

The Select DNS Inspect Map dialog appears.

Step 6   Choose the inspection map:

   - To use the default map, click **Use the default DNS inspection map** (preset_dns_map).

   - To use a DNS inspection policy map that you configured in the (Optional) Configuring a DNS Inspection Policy Map and Class Map, page 9-2, select the map name.

   - To add a new map, click **Add**. See (Optional) Configuring a DNS Inspection Policy Map and Class Map, page 9-2 for more information.

Step 7   If you use the Botnet Traffic Filter, click **Enable Botnet traffic filter DNS snooping**. Botnet Traffic Filter snooping compares the domain name with those on the dynamic database or static database, and adds the name and IP address to the Botnet Traffic Filter DNS reverse lookup cache. This cache is then used by the Botnet Traffic Filter when connections are made to the suspicious address. We suggest that you enable DNS snooping only on interfaces where external DNS requests are going. Enabling DNS snooping on all UDP DNS traffic, including that going to an internal DNS server, creates unnecessary load on the ASA. For example, if the DNS server is on the outside interface, you should enable DNS inspection with snooping for all UDP DNS traffic on the outside interface. See Enabling DNS Snooping, page 24-9.

Step 8   Click **OK** to return to the Protocol Inspections tab.

Step 9   Click **OK** to finish editing the service policy.

Step 10   Click **Apply**.

FTP Inspection

This section describes the FTP inspection engine. This section includes the following topics:

FTP Inspection Overview

The FTP application inspection inspects the FTP sessions and performs four tasks:

- Prepares dynamic secondary data connection
- Tracks the FTP command-response sequence
- Generates an audit trail
- Translates the embedded IP address

FTP application inspection prepares secondary channels for FTP data transfer. Ports for these channels are negotiated through PORT or PASV commands. The channels are allocated in response to a file upload, a file download, or a directory listing event.

> **Note**   If you disable FTP inspection engines with the **no inspect ftp** command, outbound users can start connections only in passive mode, and all inbound FTP is disabled.

Using Strict FTP

Using strict FTP increases the security of protected networks by preventing web browsers from sending embedded commands in FTP requests. To enable strict FTP, click the **Configure** button next to FTP on the Configuration > Firewall > Service Policy Rules > Edit Service Policy Rule > Rule Actions > Protocol Inspection tab.

After you enable the **strict** option on an interface, FTP inspection enforces the following behavior:

- An FTP command must be acknowledged before the ASA allows a new command.
- The ASA drops connections that send embedded commands.

- The 227 and PORT commands are checked to ensure they do not appear in an error string.

⚠ **Caution**   Using the **strict** option may cause the failure of FTP clients that are not strictly compliant with FTP RFCs.

If the **strict** option is enabled, each FTP command and response sequence is tracked for the following anomalous activity:

- Truncated command—Number of commas in the PORT and PASV reply command is checked to see if it is five. If it is not five, then the PORT command is assumed to be truncated and the TCP connection is closed.

- Incorrect command—Checks the FTP command to see if it ends with <CR><LF> characters, as required by the RFC. If it does not, the connection is closed.

- Size of RETR and STOR commands—These are checked against a fixed constant. If the size is greater, then an error message is logged and the connection is closed.

- Command spoofing—The PORT command should always be sent from the client. The TCP connection is denied if a PORT command is sent from the server.

- Reply spoofing—PASV reply command (227) should always be sent from the server. The TCP connection is denied if a PASV reply command is sent from the client. This prevents the security hole when the user executes "227 xxxxx a1, a2, a3, a4, p1, p2."

- TCP stream editing—The ASA closes the connection if it detects TCP stream editing.

- Invalid port negotiation—The negotiated dynamic port value is checked to see if it is less than 1024. As port numbers in the range from 1 to 1024 are reserved for well-known connections, if the negotiated port falls in this range, then the TCP connection is freed.

- Command pipelining—The number of characters present after the port numbers in the PORT and PASV reply command is cross checked with a constant value of 8. If it is more than 8, then the TCP connection is closed.

- The ASA replaces the FTP server response to the SYST command with a series of Xs. to prevent the server from revealing its system type to FTP clients. To override this default behavior, use the **no mask-syst-reply** command in the FTP map.

Select FTP Map

The Select FTP Map dialog box is accessible as follows:

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select FTP Map

The Select FTP Map dialog box lets you enable strict FTP application inspection, select an FTP map, or create a new FTP map. An FTP map lets you change the configuration values used for FTP application inspection.The Select FTP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- FTP Strict (prevent web browsers from sending embedded commands in FTP requests)**—**Enables strict FTP application inspection, which causes the ASA to drop the connection when an embedded command is included in an FTP request.

- Use the default FTP inspection map—Specifies to use the default FTP map.

- Select an FTP map for fine control over inspection—Lets you select a defined application inspection map or add a new one.

- Add—Opens the Add Policy Map dialog box for the inspection.

FTP Class Map

The FTP Class Map dialog box is accessible as follows:

Configuration > Global Objects > Class Maps > FTP

The FTP Class Map pane lets you configure FTP class maps for FTP inspection.

An inspection class map matches application traffic with criteria specific to the application. You then identify the class map in the inspect map and enable actions. The difference between creating a class map and defining the traffic match directly in the inspect map is that you can create more complex match criteria and you can reuse class maps. The applications that support inspection class maps are DNS, FTP, H.323, HTTP, IM, and SIP.

Fields

- Name—Shows the FTP class map name.
- Match Conditions—Shows the type, match criterion, and value in the class map.
    - Match Type—Shows the match type, which can be a positive or negative match.
    - Criterion—Shows the criterion of the FTP class map.
    - Value—Shows the value to match in the FTP class map.
- Description—Shows the description of the class map.
- Add—Adds an FTP class map.
- Edit—Edits an FTP class map.
- Delete—Deletes an FTP class map.

Add/Edit FTP Traffic Class Map

The Add/Edit FTP Traffic Class Map dialog box is accessible as follows:

Configuration > Global Objects > Class Maps > FTP > Add/Edit FTP Traffic Class Map

The Add/Edit FTP Traffic Class Map dialog box lets you define a FTP class map.

Fields

- Name—Enter the name of the FTP class map, up to 40 characters in length.
- Description—Enter the description of the FTP class map.
- Add—Adds an FTP class map.
- Edit—Edits an FTP class map.
- Delete—Deletes an FTP class map.

Add/Edit FTP Match Criterion

The Add/Edit FTP Match Criterion dialog box is accessible as follows:

Configuration > Global Objects > Class Maps > FTP > Add/Edit FTP Traffic Class Map > Add/Edit FTP Match Criterion

The Add/Edit FTP Match Criterion dialog box lets you define the match criterion and value for the FTP class map.

Fields

- Match Type—Specifies whether the class map includes traffic that matches the criterion, or traffic that does not match the criterion.

    For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map.

- Criterion—Specifies which criterion of FTP traffic to match.

    – Request-Command—Match an FTP request command.

    – File Name—Match a filename for FTP transfer.

    – File Type—Match a file type for FTP transfer.

    – Server—Match an FTP server.

    – User Name—Match an FTP user.

- Request-Command Criterion Values—Specifies the value details for the FTP request command match.

    – Request Command—Lets you select one or more request commands to match.

        APPE—Append to a file.

        CDUP—Change to the parent of the current directory.

        DELE—Delete a file at the server site.

        GET—FTP client command for the retr (retrieve a file) command.

        HELP—Help information from the server.

        MKD—Create a directory.

        PUT—FTP client command for the stor (store a file) command.

        RMD—Remove a directory.

        RNFR—Rename from.

        RNTO—Rename to.

        SITE—Specify a server specific command.

        STOU—Store a file with a unique name.

- File Name Criterion Values—Specifies to match on the FTP transfer filename.

    – Regular Expression—Lists the defined regular expressions to match.

    – Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

    – Regular Expression Class—Lists the defined regular expression classes to match.

    – Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- File Type Criterion Values—Specifies to match on the FTP transfer file type.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Server Criterion Values—Specifies to match on the FTP server.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- User Name Criterion Values—Specifies to match on the FTP user.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

FTP Inspect Map

The FTP Inspect Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > FTP

The FTP pane lets you view previously configured FTP application inspection maps. An FTP map lets you change the default configuration values used for FTP application inspection.

FTP command filtering and security checks are provided using strict FTP inspection for improved security and control. Protocol conformance includes packet length checks, delimiters and packet format checks, command terminator checks, and command validation.

Blocking FTP based on user values is also supported so that it is possible for FTP sites to post files for download, but restrict access to certain users. You can block FTP connections based on file type, server name, and other attributes. System message logs are generated if an FTP connection is denied after inspection.

Fields

- FTP Inspect Maps—Table that lists the defined FTP inspect maps.
- Add—Configures a new FTP inspect map. To edit an FTP inspect map, choose the FTP entry in the FTP Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the FTP Inspect Maps table.
- Security Level—Select the security level (medium or low).

> – Low

  Mask Banner Disabled

  Mask Reply Disabled

> – Medium—Default.

  Mask Banner Enabled

  Mask Reply Enabled

> – File Type Filtering—Opens the Type Filtering dialog box to configure file type filters.

> – Customize—Opens the Add/Edit FTP Policy Map dialog box for additional settings.

> – Default Level—Sets the security level back to the default level of Medium.

File Type Filtering

The File Type Filtering dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > FTP > MIME File Type Filtering

The File Type Filtering dialog box lets you configure the settings for a file type filter.

Fields

- Match Type—Shows the match type, which can be a positive or negative match.

- Criterion—Shows the criterion of the inspection.

- Value—Shows the value to match in the inspection.

- Action—Shows the action if the match condition is met.

- Log—Shows the log state.

- Add—Opens the Add File Type Filter dialog box to add a file type filter.

- Edit—Opens the Edit File Type Filter dialog box to edit a file type filter.

- Delete—Deletes a file type filter.

- Move Up—Moves an entry up in the list.

- Move Down—Moves an entry down in the list.

Add/Edit FTP Policy Map (Security Level)

The Add/Edit FTP Policy Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > FTP > FTP Inspect Map > Basic View

The Add/Edit FTP Policy Map pane lets you configure the security level and additional settings for FTP application inspection maps.

Fields

- Name—When adding an FTP map, enter the name of the FTP map. When editing an FTP map, the name of the previously configured FTP map is shown.

- Description—Enter the description of the FTP map, up to 200 characters in length.

- Security Level—Select the security level (medium or low).

        – Low

          Mask Banner Disabled

          Mask Reply Disabled

        – Medium—Default.

          Mask Banner Enabled

          Mask Reply Enabled

        – File Type Filtering—Opens the Type Filtering dialog box to configure file type filters.

        – Default Level—Sets the security level back to the default level of Medium.

- Details—Shows the Parameters and Inspections tabs to configure additional settings.

Add/Edit FTP Policy Map (Details)

The Add/Edit FTP Policy Map (Details) dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > FTP > FTP Inspect Map > Advanced View

The Add/Edit FTP Policy Map pane lets you configure the security level and additional settings for FTP application inspection maps.

Fields

- Name—When adding an FTP map, enter the name of the FTP map. When editing an FTP map, the name of the previously configured FTP map is shown.

- Description—Enter the description of the FTP map, up to 200 characters in length.

- Security Level—Shows the security level and file type filtering settings to configure.

- Parameters—Tab that lets you configure the parameters for the FTP inspect map.

        – Mask greeting banner from the server—Masks the greeting banner from the FTP server to prevent the client from discovering server information.

        – Mask reply to SYST command—Masks the reply to the syst command to prevent the client from discovering server information.

- Inspections—Tab that shows you the FTP inspection configuration and lets you add or edit.

        – Match Type—Shows the match type, which can be a positive or negative match.

        – Criterion—Shows the criterion of the FTP inspection.

        – Value—Shows the value to match in the FTP inspection.

        – Action—Shows the action if the match condition is met.

        – Log—Shows the log state.

        – Add—Opens the Add FTP Inspect dialog box to add an FTP inspection.

        – Edit—Opens the Edit FTP Inspect dialog box to edit an FTP inspection.

        – Delete—Deletes an FTP inspection.

        – Move Up—Moves an inspection up in the list.

        – Move Down—Moves an inspection down in the list.

Add/Edit FTP Map

The Add/Edit FTP Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > FTP > FTP Inspect Map > Advanced View > Add/Edit FTP Inspect

The Add/Edit FTP Inspect dialog box lets you define the match criterion and value for the FTP inspect map.

Fields

- Single Match—Specifies that the FTP inspect has only one match statement.
- Match Type—Specifies whether traffic should match or not match the values.

  For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map.

- Criterion—Specifies which criterion of FTP traffic to match.
    - Request Command—Match an FTP request command.
    - File Name—Match a filename for FTP transfer.
    - File Type—Match a file type for FTP transfer.
    - Server—Match an FTP server.
    - User Name—Match an FTP user.
- Request Command Criterion Values—Specifies the value details for FTP request command match.
    - Request Command:

      APPE—Command that appends to a file.

      CDUP—Command that changes to the parent directory of the current working directory.

      DELE—Command that deletes a file.

      GET—Command that gets a file.

      HELP—Command that provides help information.

      MKD—Command that creates a directory.

      PUT—Command that sends a file.

      RMD—Command that deletes a directory.

      RNFR—Command that specifies rename-from filename.

      RNTO—Command that specifies rename-to filename.

      SITE—Commands that are specific to the server system. Usually used for remote administration.

      STOU—Command that stores a file using a unique filename.
- File Name Criterion Values—Specifies the value details for FTP filename match.
    - Regular Expression—Lists the defined regular expressions to match.
    - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
    - Regular Expression Class—Lists the defined regular expression classes to match.

- Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- File Type Criterion Values—Specifies the value details for FTP file type match.

    - Regular Expression—Lists the defined regular expressions to match.

    - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

    - Regular Expression Class—Lists the defined regular expression classes to match.

    - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Server Criterion Values—Specifies the value details for FTP server match.

    - Regular Expression—Lists the defined regular expressions to match.

    - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

    - Regular Expression Class—Lists the defined regular expression classes to match.

    - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- User Name Criterion Values—Specifies the value details for FTP user name match.

    - Regular Expression—Lists the defined regular expressions to match.

    - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

    - Regular Expression Class—Lists the defined regular expression classes to match.

    - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Multiple Matches—Specifies multiple matches for the FTP inspection.

    - FTP Traffic Class—Specifies the FTP traffic class match.

    - Manage—Opens the Manage FTP Class Maps dialog box to add, edit, or delete FTP Class Maps.

- Action—Reset.

- Log—Enable or disable.

Verifying and Monitoring FTP Inspection

FTP application inspection generates the following log messages:

- An Audit record 303002 is generated for each file that is retrieved or uploaded.

- The FTP command is checked to see if it is RETR or STOR and the retrieve and store commands are logged.

- The username is obtained by looking up a table providing the IP address.

- The username, source IP address, destination IP address, NAT address, and the file operation are logged.

- Audit record 201005 is generated if the secondary dynamic channel preparation failed due to memory shortage.

In conjunction with NAT, the FTP application inspection translates the IP address within the application payload. This is described in detail in RFC 959.

HTTP Inspection

This section describes the HTTP inspection engine. This section includes the following topics:

HTTP Inspection Overview

Use the HTTP inspection engine to protect against specific attacks and other threats that are associated with HTTP traffic. The latter two features are configured in conjunction with Filter rules.

The enhanced HTTP inspection feature, which is also known as an application firewall and is available when you configure an HTTP map, can help prevent attackers from using HTTP messages for circumventing network security policy. It verifies the following for all HTTP messages:

- Conformance to RFC 2616
- Use of RFC-defined methods only.
- Compliance with the additional criteria.

Select HTTP Map

The Select HTTP Map dialog box is accessible as follows:

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select HTTP Map

The Select HTTP Map dialog box lets you select or create a new HTTP map. An HTTP map lets you change the configuration values used for HTTP application inspection. The Select HTTP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default HTTP inspection map—Specifies to use the default HTTP map.
- Select an HTTP map for fine control over inspection—Lets you select a defined application inspection map or add a new one.

- Add—Opens the Add Policy Map dialog box for the inspection.

HTTP Class Map

The HTTP Class Map dialog box is accessible as follows:

Configuration > Global Objects > Class Maps > HTTP

The HTTP Class Map pane lets you configure HTTP class maps for HTTP inspection.

An inspection class map matches application traffic with criteria specific to the application. You then identify the class map in the inspect map and enable actions. The difference between creating a class map and defining the traffic match directly in the inspect map is that you can create more complex match criteria and you can reuse class maps. The applications that support inspection class maps are DNS, FTP, H.323, HTTP, IM, and SIP.

Fields

- Name—Shows the HTTP class map name.
- Match Conditions—Shows the type, match criterion, and value in the class map.
    - Match Type—Shows the match type, which can be a positive or negative match.
    - Criterion—Shows the criterion of the HTTP class map.
    - Value—Shows the value to match in the HTTP class map.
- Description—Shows the description of the class map.
- Add—Adds an HTTP class map.
- Edit—Edits an HTTP class map.
- Delete—Deletes an HTTP class map.

Add/Edit HTTP Traffic Class Map

The Add/Edit HTTP Traffic Class Map dialog box is accessible as follows:

Configuration > Global Objects > Class Maps > HTTP > Add/Edit HTTP Traffic Class Map

The Add/Edit HTTP Traffic Class Map dialog box lets you define a HTTP class map.

Fields

- Name—Enter the name of the HTTP class map, up to 40 characters in length.
- Description—Enter the description of the HTTP class map.
- Add—Adds an HTTP class map.
- Edit—Edits an HTTP class map.
- Delete—Deletes an HTTP class map.

Add/Edit HTTP Match Criterion

The Add/Edit HTTP Match Criterion dialog box is accessible as follows:

Configuration > Global Objects > Class Maps > HTTP > Add/Edit HTTP Traffic Class Map > Add/Edit HTTP Match Criterion

The Add/Edit HTTP Match Criterion dialog box lets you define the match criterion and value for the HTTP class map.

Fields

- Match Type—Specifies whether the class map includes traffic that matches the criterion, or traffic that does not match the criterion.

  For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map.

- Criterion—Specifies which criterion of HTTP traffic to match.

  - Request/Response Content Type Mismatch—Specifies that the content type in the response must match one of the MIME types in the accept field of the request.

  - Request Arguments—Applies the regular expression match to the arguments of the request.

    Regular Expression—Lists the defined regular expressions to match.

    Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

    Regular Expression Class—Lists the defined regular expression classes to match.

    Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

  - Request Body Length—Applies the regular expression match to the body of the request with field length greater than the bytes specified.

    Greater Than Length—Enter a field length value in bytes that request field lengths will be matched against.

  - Request Body—Applies the regular expression match to the body of the request.

    Regular Expression—Lists the defined regular expressions to match.

    Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

    Regular Expression Class—Lists the defined regular expression classes to match.

    Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

  - Request Header Field Count—Applies the regular expression match to the header of the request with a maximum number of header fields.

    Predefined—Specifies the request header fields: accept, accept-charset, accept-encoding, accept-language, allow, authorization, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, cookie, date, expect, expires, from, host, if-match, if-modified-since, if-none-match, if-range, if-unmodified-since, last-modified, max-forwards, pragma, proxy-authorization, range, referer, te, trailer, transfer-encoding, upgrade, user-agent, via, warning.

    Regular Expression—Lists the defined regular expressions to match.

    Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

    Greater Than Count—Enter the maximum number of header fields.

- Request Header Field Length—Applies the regular expression match to the header of the request with field length greater than the bytes specified.

  Predefined—Specifies the request header fields: accept, accept-charset, accept-encoding, accept-language, allow, authorization, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, cookie, date, expect, expires, from, host, if-match, if-modified-since, if-none-match, if-range, if-unmodified-since, last-modified, max-forwards, pragma, proxy-authorization, range, referer, te, trailer, transfer-encoding, upgrade, user-agent, via, warning.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Greater Than Length—Enter a field length value in bytes that request field lengths will be matched against.

- Request Header Field—Applies the regular expression match to the header of the request.

  Predefined—Specifies the request header fields: accept, accept-charset, accept-encoding, accept-language, allow, authorization, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, cookie, date, expect, expires, from, host, if-match, if-modified-since, if-none-match, if-range, if-unmodified-since, last-modified, max-forwards, pragma, proxy-authorization, range, referer, te, trailer, transfer-encoding, upgrade, user-agent, via, warning.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Regular Expression Class—Lists the defined regular expression classes to match.

  Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Request Header Count—Applies the regular expression match to the header of the request with a maximum number of headers.

  Greater Than Count—Enter the maximum number of headers.

- Request Header Length—Applies the regular expression match to the header of the request with length greater than the bytes specified.

  Greater Than Length—Enter a header length value in bytes.

- Request Header non-ASCII—Matches non-ASCII characters in the header of the request.

- Request Method—Applies the regular expression match to the method of the request.

  Method—Specifies to match on a request method: bcopy, bdelete, bmove, bpropfind, bproppatch, connect, copy, delete, edit, get, getattribute, getattributenames, getproperties, head, index, lock, mkcol, mkdir, move, notify, options, poll, post, propfind, proppatch, put, revadd, revlabel, revlog, revnum, save, search, setattribute, startrev, stoprev, subscribe, trace, unedit, unlock, unsubscribe.

  Regular Expression—Specifies to match on a regular expression.

  Regular Expression—Lists the defined regular expressions to match.

  Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  Regular Expression Class—Lists the defined regular expression classes to match.

        Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

  – Request URI Length—Applies the regular expression match to the URI of the request with length greater than the bytes specified.

        Greater Than Length—Enter a URI length value in bytes.

  – Request URI—Applies the regular expression match to the URI of the request.

        Regular Expression—Lists the defined regular expressions to match.

        Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

        Regular Expression Class—Lists the defined regular expression classes to match.

        Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

  – Response Body—Applies the regex match to the body of the response.

        ActiveX—Specifies to match on ActiveX.

        Java Applet—Specifies to match on a Java Applet.

        Regular Expression—Specifies to match on a regular expression.

        Regular Expression—Lists the defined regular expressions to match.

        Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

        Regular Expression Class—Lists the defined regular expression classes to match.

        Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

  – Response Body Length—Applies the regular expression match to the body of the response with field length greater than the bytes specified.

        Greater Than Length—Enter a field length value in bytes that response field lengths will be matched against.

  – Response Header Field Count—Applies the regular expression match to the header of the response with a maximum number of header fields.

        Predefined—Specifies the response header fields: accept-ranges, age, allow, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, date, etag, expires, last-modified, location, pragma, proxy-authenticate, retry-after, server, set-cookie, trailer, transfer-encoding, upgrade, vary, via, warning, www-authenticate.

        Regular Expression—Lists the defined regular expressions to match.

        Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

        Greater Than Count—Enter the maximum number of header fields.

  – Response Header Field Length—Applies the regular expression match to the header of the response with field length greater than the bytes specified.

Predefined—Specifies the response header fields: accept-ranges, age, allow, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, date, etag, expires, last-modified, location, pragma, proxy-authenticate, retry-after, server, set-cookie, trailer, transfer-encoding, upgrade, vary, via, warning, www-authenticate.

Regular Expression—Lists the defined regular expressions to match.

Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

Greater Than Length—Enter a field length value in bytes that response field lengths will be matched against.

- Response Header Field—Applies the regular expression match to the header of the response.

Predefined—Specifies the response header fields: accept-ranges, age, allow, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, date, etag, expires, last-modified, location, pragma, proxy-authenticate, retry-after, server, set-cookie, trailer, transfer-encoding, upgrade, vary, via, warning, www-authenticate.

Regular Expression—Lists the defined regular expressions to match.

Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

Regular Expression Class—Lists the defined regular expression classes to match.

Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Response Header Count—Applies the regular expression match to the header of the response with a maximum number of headers.

Greater Than Count—Enter the maximum number of headers.

- Response Header Length—Applies the regular expression match to the header of the response with length greater than the bytes specified.

Greater Than Length—Enter a header length value in bytes.

- Response Header non-ASCII—Matches non-ASCII characters in the header of the response.

- Response Status Line—Applies the regular expression match to the status line.

Regular Expression—Lists the defined regular expressions to match.

Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

Regular Expression Class—Lists the defined regular expression classes to match.

Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

HTTP Inspect Map

The HTTP Inspect Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > HTTP

The HTTP pane lets you view previously configured HTTP application inspection maps. An HTTP map lets you change the default configuration values used for HTTP application inspection.

HTTP application inspection scans HTTP headers and body, and performs various checks on the data. These checks prevent various HTTP constructs, content types, and tunneling and messaging protocols from traversing the security appliance.

HTTP application inspection can block tunneled applications and non-ASCII characters in HTTP requests and responses, preventing malicious content from reaching the web server. Size limiting of various elements in HTTP request and response headers, URL blocking, and HTTP server header type spoofing are also supported.

Fields

- HTTP Inspect Maps—Table that lists the defined HTTP inspect maps.

- Add—Configures a new HTTP inspect map. To edit an HTTP inspect map, choose the HTTP entry in the HTTP Inspect Maps table and click **Customize**.

- Delete—Deletes the inspect map selected in the HTTP Inspect Maps table.

- Security Level—Select the security level (low, medium, or high).

    - Low—Default.

        Protocol violation action: Drop connection

        Drop connections for unsafe methods: Disabled

        Drop connections for requests with non-ASCII headers: Disabled

        URI filtering: Not configured

        Advanced inspections: Not configured

    - Medium

        Protocol violation action: Drop connection

        Drop connections for unsafe methods: Allow only GET, HEAD, and POST

        Drop connections for requests with non-ASCII headers: Disabled

        URI filtering: Not configured

        Advanced inspections: Not configured

    - High

        Protocol violation action: Drop connection and log

        Drop connections for unsafe methods: Allow only GET and HEAD.

        Drop connections for requests with non-ASCII headers: Enabled

        URI filtering: Not configured

        Advanced inspections: Not configured

    - URI Filtering—Opens the URI Filtering dialog box to configure URI filters.

    - Customize—Opens the Edit HTTP Policy Map dialog box for additional settings.

    - Default Level—Sets the security level back to the default level of Medium.

URI Filtering

The URI Filtering dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > HTTP > URI Filtering

The URI Filtering dialog box lets you configure the settings for an URI filter.

Fields

- Match Type—Shows the match type, which can be a positive or negative match.
- Criterion—Shows the criterion of the inspection.
- Value—Shows the value to match in the inspection.
- Action—Shows the action if the match condition is met.
- Log—Shows the log state.
- Add—Opens the Add URI Filtering dialog box to add a URI filter.
- Edit—Opens the Edit URI Filtering dialog box to edit a URI filter.
- Delete—Deletes an URI filter.
- Move Up—Moves an entry up in the list.
- Move Down—Moves an entry down in the list.

Add/Edit HTTP Policy Map (Security Level)

The Add/Edit HTTP Policy Map (Security Level) dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > HTTP > HTTP Inspect Map > Basic View

The Add/Edit HTTP Policy Map pane lets you configure the security level and additional settings for HTTP application inspection maps.

Fields

- Name—When adding an HTTP map, enter the name of the HTTP map. When editing an HTTP map, the name of the previously configured HTTP map is shown.
- Description—Enter the description of the HTTP map, up to 200 characters in length.
- Security Level—Select the security level (low, medium, or high).
  - Low—Default.

    Protocol violation action: Drop connection

    Drop connections for unsafe methods: Disabled

    Drop connections for requests with non-ASCII headers: Disabled

    URI filtering: Not configured

    Advanced inspections: Not configured
  - Medium

    Protocol violation action: Drop connection

    Drop connections for unsafe methods: Allow only GET, HEAD, and POST

    Drop connections for requests with non-ASCII headers: Disabled

    URI filtering: Not configured

    Advanced inspections: Not configured
  - High

    Protocol violation action: Drop connection and log

Drop connections for unsafe methods: Allow only GET and HEAD.

Drop connections for requests with non-ASCII headers: Enabled

URI filtering: Not configured

Advanced inspections: Not configured

- – URI Filtering—Opens the URI Filtering dialog box which lets you configure the settings for an URI filter.

- – Default Level—Sets the security level back to the default.

- • Details—Shows the Parameters and Inspections tabs to configure additional settings.

Add/Edit HTTP Policy Map (Details)

The Add/Edit HTTP Policy Map (Details) dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > HTTP > HTTP Inspect Map > Advanced View

The Add/Edit HTTP Policy Map pane lets you configure the security level and additional settings for HTTP application inspection maps.

Fields

- • Name—When adding an HTTP map, enter the name of the HTTP map. When editing an HTTP map, the name of the previously configured HTTP map is shown.

- • Description—Enter the description of the HTTP map, up to 200 characters in length.

- • Security Level—Shows the security level and URI filtering settings to configure.

- • Parameters—Tab that lets you configure the parameters for the HTTP inspect map.

   - – Check for protocol violations—Checks for HTTP protocol violations.

      Action—Drop Connection, Reset, Log.

      Log—Enable or disable.

   - – Spoof server string—Replaces the server HTTP header value with the specified string.

      Spoof String—Enter a string to substitute for the server header field. Maximum is 82 characters.

   - – Body Match Maximum—The maximum number of characters in the body of an HTTP message that should be searched in a body match. Default is 200 bytes. A large number will have a significant impact on performance.

- • Inspections—Tab that shows you the HTTP inspection configuration and lets you add or edit.

   - – Match Type—Shows the match type, which can be a positive or negative match.

   - – Criterion—Shows the criterion of the HTTP inspection.

   - – Value—Shows the value to match in the HTTP inspection.

   - – Action—Shows the action if the match condition is met.

   - – Log—Shows the log state.

   - – Add—Opens the Add HTTP Inspect dialog box to add an HTTP inspection.

   - – Edit—Opens the Edit HTTP Inspect dialog box to edit an HTTP inspection.

   - – Delete—Deletes an HTTP inspection.

   - – Move Up—Moves an inspection up in the list.

– Move Down—Moves an inspection down in the list.

Add/Edit HTTP Map

The Add/Edit HTTP Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > HTTP > HTTP Inspect Map > Advanced View > Add/Edit HTTP Inspect

The Add/Edit HTTP Inspect dialog box lets you define the match criterion and value for the HTTP inspect map.

Fields

- Single Match—Specifies that the HTTP inspect has only one match statement.

- Match Type—Specifies whether traffic should match or not match the values.

  For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map.

- Criterion—Specifies which criterion of HTTP traffic to match.

  – Request/Response Content Type Mismatch—Specifies that the content type in the response must match one of the MIME types in the accept field of the request.

  – Request Arguments—Applies the regular expression match to the arguments of the request.

    Regular Expression—Lists the defined regular expressions to match.

    Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

    Regular Expression Class—Lists the defined regular expression classes to match.

    Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

  – Request Body Length—Applies the regular expression match to the body of the request with field length greater than the bytes specified.

    Greater Than Length—Enter a field length value in bytes that request field lengths will be matched against.

  – Request Body—Applies the regular expression match to the body of the request.

    Regular Expression—Lists the defined regular expressions to match.

    Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

    Regular Expression Class—Lists the defined regular expression classes to match.

    Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

  – Request Header Field Count—Applies the regular expression match to the header of the request with a maximum number of header fields.

    Predefined—Specifies the request header fields: accept, accept-charset, accept-encoding, accept-language, allow, authorization, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type,

cookie, date, expect, expires, from, host, if-match, if-modified-since, if-none-match, if-range, if-unmodified-since, last-modified, max-forwards, pragma, proxy-authorization, range, referer, te, trailer, transfer-encoding, upgrade, user-agent, via, warning.

Regular Expression—Lists the defined regular expressions to match.

Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

Greater Than Count—Enter the maximum number of header fields.

– Request Header Field Length—Applies the regular expression match to the header of the request with field length greater than the bytes specified.

Predefined—Specifies the request header fields: accept, accept-charset, accept-encoding, accept-language, allow, authorization, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, cookie, date, expect, expires, from, host, if-match, if-modified-since, if-none-match, if-range, if-unmodified-since, last-modified, max-forwards, pragma, proxy-authorization, range, referer, te, trailer, transfer-encoding, upgrade, user-agent, via, warning.

Regular Expression—Lists the defined regular expressions to match.

Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

Greater Than Length—Enter a field length value in bytes that request field lengths will be matched against.

– Request Header Field—Applies the regular expression match to the header of the request.

Predefined—Specifies the request header fields: accept, accept-charset, accept-encoding, accept-language, allow, authorization, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, cookie, date, expect, expires, from, host, if-match, if-modified-since, if-none-match, if-range, if-unmodified-since, last-modified, max-forwards, pragma, proxy-authorization, range, referer, te, trailer, transfer-encoding, upgrade, user-agent, via, warning.

Regular Expression—Lists the defined regular expressions to match.

Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

Regular Expression Class—Lists the defined regular expression classes to match.

Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

– Request Header Count—Applies the regular expression match to the header of the request with a maximum number of headers.

Greater Than Count—Enter the maximum number of headers.

– Request Header Length—Applies the regular expression match to the header of the request with length greater than the bytes specified.

Greater Than Length—Enter a header length value in bytes.

– Request Header non-ASCII—Matches non-ASCII characters in the header of the request.

– Request Method—Applies the regular expression match to the method of the request.

Method—Specifies to match on a request method: bcopy, bdelete, bmove, bpropfind, bproppatch, connect, copy, delete, edit, get, getattribute, getattributenames, getproperties, head, index, lock, mkcol, mkdir, move, notify, options, poll, post, propfind, proppatch, put, revadd, revlabel, revlog, revnum, save, search, setattribute, startrev, stoprev, subscribe, trace, unedit, unlock, unsubscribe.

Regular Expression—Specifies to match on a regular expression.

Regular Expression—Lists the defined regular expressions to match.

Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

Regular Expression Class—Lists the defined regular expression classes to match.

Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

– Request URI Length—Applies the regular expression match to the URI of the request with length greater than the bytes specified.

Greater Than Length—Enter a URI length value in bytes.

– Request URI—Applies the regular expression match to the URI of the request.

Regular Expression—Lists the defined regular expressions to match.

Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

Regular Expression Class—Lists the defined regular expression classes to match.

Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

– Response Body—Applies the regex match to the body of the response.

ActiveX—Specifies to match on ActiveX.

Java Applet—Specifies to match on a Java Applet.

Regular Expression—Specifies to match on a regular expression.

Regular Expression—Lists the defined regular expressions to match.

Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

Regular Expression Class—Lists the defined regular expression classes to match.

Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

– Response Body Length—Applies the regular expression match to the body of the response with field length greater than the bytes specified.

Greater Than Length—Enter a field length value in bytes that response field lengths will be matched against.

– Response Header Field Count—Applies the regular expression match to the header of the response with a maximum number of header fields.

Predefined—Specifies the response header fields: accept-ranges, age, allow, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, date, etag, expires, last-modified, location, pragma, proxy-authenticate, retry-after, server, set-cookie, trailer, transfer-encoding, upgrade, vary, via, warning, www-authenticate.

Regular Expression—Lists the defined regular expressions to match.

Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

Greater Than Count—Enter the maximum number of header fields.

– Response Header Field Length—Applies the regular expression match to the header of the response with field length greater than the bytes specified.

Predefined—Specifies the response header fields: accept-ranges, age, allow, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, date, etag, expires, last-modified, location, pragma, proxy-authenticate, retry-after, server, set-cookie, trailer, transfer-encoding, upgrade, vary, via, warning, www-authenticate.

Regular Expression—Lists the defined regular expressions to match.

Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

Greater Than Length—Enter a field length value in bytes that response field lengths will be matched against.

– Response Header Field—Applies the regular expression match to the header of the response.

Predefined—Specifies the response header fields: accept-ranges, age, allow, cache-control, connection, content-encoding, content-language, content-length, content-location, content-md5, content-range, content-type, date, etag, expires, last-modified, location, pragma, proxy-authenticate, retry-after, server, set-cookie, trailer, transfer-encoding, upgrade, vary, via, warning, www-authenticate.

Regular Expression—Lists the defined regular expressions to match.

Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

Regular Expression Class—Lists the defined regular expression classes to match.

Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

– Response Header Count—Applies the regular expression match to the header of the response with a maximum number of headers.

Greater Than Count—Enter the maximum number of headers.

– Response Header Length—Applies the regular expression match to the header of the response with length greater than the bytes specified.

Greater Than Length—Enter a header length value in bytes.

– Response Header non-ASCII—Matches non-ASCII characters in the header of the response.

– Response Status Line—Applies the regular expression match to the status line.

Regular Expression—Lists the defined regular expressions to match.

Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

Regular Expression Class—Lists the defined regular expression classes to match.

Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

• Multiple Matches—Specifies multiple matches for the HTTP inspection.

- **H323 Traffic Class**—Specifies the HTTP traffic class match.

- **Manage**—Opens the Manage HTTP Class Maps dialog box to add, edit, or delete HTTP Class Maps.

- **Action**—Drop connection, reset, or log.

- **Log**—Enable or disable.

ICMP Inspection

The ICMP inspection engine allows ICMP traffic to have a "session" so it can be inspected like TCP and UDP traffic. Without the ICMP inspection engine, we recommend that you do not allow ICMP through the ASA in an ACL. Without stateful inspection, ICMP can be used to attack your network. The ICMP inspection engine ensures that there is only one response for each request, and that the sequence number is correct.

ICMP Error Inspection

When this feature is enabled, the ASA creates translation sessions for intermediate hops that send ICMP error messages, based on the NAT configuration. The ASA overwrites the packet with the translated IP addresses.

When disabled, the ASA does not create translation sessions for intermediate nodes that generate ICMP error messages. ICMP error messages generated by the intermediate nodes between the inside host and the ASA reach the outside host without consuming any additional NAT resource. This is undesirable when an outside host uses the traceroute command to trace the hops to the destination on the inside of the ASA. When the ASA does not translate the intermediate hops, all the intermediate hops appear with the mapped destination IP address.

The ICMP payload is scanned to retrieve the five-tuple from the original packet. Using the retrieved five-tuple, a lookup is performed to determine the original address of the client. The ICMP error inspection engine makes the following changes to the ICMP packet:

- In the IP Header, the mapped IP is changed to the real IP (Destination Address) and the IP checksum is modified.

- In the ICMP Header, the ICMP checksum is modified due to the changes in the ICMP packet.

- In the Payload, the following changes are made:

    - Original packet mapped IP is changed to the real IP

    - Original packet mapped port is changed to the real Port

    - Original packet IP checksum is recalculated

Instant Messaging Inspection

This section describes the IM inspection engine. This section includes the following topics:

IM Inspection Overview

The IM inspect engine lets you apply fine grained controls on the IM application to control the network usage and stop leakage of confidential data, propagation of worms, and other threats to the corporate network.

Adding a Class Map for IM Inspection

Use the Add Service Policy Rule Wizard - Rule Actions dialog box to configure IP Options inspection.

This wizard is available from the Configuration > Firewall > Service Policy Rules > Add > Add Service Policy Rule Wizard - Rule Actions dialog box.

Step 1   Choose **Configuration > Firewall > Objects > Class Maps > Instant Messaging (IM)**. The table displaying the configured class maps for Instant Messaging Inspection appears.

Step 2   To add a new class map, click **Add**. The Add Instant Messaging (IM) Traffic Class Map dialog box appears.

Step 3   Enter a name for the class map.

Step 4   (Optional) Enter a description for the class map. The description can contain up to 200 characters.

Step 5   In the Match Option field, click an option for the class map:

- Match All—Specifies that traffic must match all criteria to match the class map. By default, the Match All option is selected.

- Match Any—Specifies that the traffic matches the class map if it matches at least one of the criteria.

Step 6   Click **Add** to add a match criteria for the class map. The Add Instant Messaging (IM) Match Criterion dialog box appears.

Step 7   In the Match Type field, click the Match or No Match radio button.

Step 8   In the Criterion drop-down list, select one of the following options and specify the criteria value. Depending on which option you select, the Value fields dynamically refresh to display the appropriate values for that criteria.

- Protocol—Select to match traffic of a specific IM protocol, such as Yahoo Messenger or MSN Messenger.

- Service—Select to match a specific IM service, such as chat, file-transfer, webcam, voice-chat, conference, or games.

- Version—Select to match the version of the IM message. In the Value fields, click the **Regular Expression** or **Regular Expression Class** option and select an expression from the drop-down list.

- Client Login Name—Select to match the source login name of the IM message. In the Value fields, click the **Regular Expression** or **Regular Expression Class** option and select an expression from the drop-down list.

- Client Peer Login Name—Select to match the destination login name of the IM message. In the Value fields, click the **Regular Expression** or **Regular Expression Class** option and select an expression from the drop-down list.

- Source IP Address—Select to match the source IP address of the IM message. In the Value fields, enter the IP address and netmask of the message source.

- Destination IP Address—Select to match the destination IP address of the IM message. In the Value fields, enter the IP address and netmask of the message destination.

- Filename—Select to match the filename of the IM message. In the Value fields, click the **Regular Expression** or **Regular Expression Class** option and select an expression from the drop-down list.

Step 9   Click **OK** to save the criteria. The Add Instant Messaging (IM) Match Criterion dialog box closes and the criteria appears in the Match Criterion table.

Step 10   Click **OK** to save the class map.

Select IM Map

The Select IM Map dialog box is accessible as follows:

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select IM Map

The Select IM Map dialog box lets you select or create a new IM map. An IM map lets you change the configuration values used for IM application inspection. The Select IM Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Add—Opens the Add Policy Map dialog box for the inspection.

IP Options Inspection

This section describes the IP Options inspection engine. This section includes the following topics:

- IP Options Inspection Overview, page 9-41
- Configuring IP Options Inspection, page 9-42
- Select IP Options Inspect Map, page 9-43
- IP Options Inspect Map, page 9-44
- Add/Edit IP Options Inspect Map, page 9-44

IP Options Inspection Overview

Each IP packet contains an IP header with the Options field. The Options field, commonly referred to as IP Options, provide for control functions that are required in some situations but unnecessary for most common communications. In particular, IP Options include provisions for time stamps, security, and special routing. Use of IP Options is optional, and the field can contain zero, one, or more options.

You can configure IP Options inspection to control which IP packets with specific IP options are allowed through the ASA. Configuring this inspection instructs the ASA to allow a packet to pass or to clear the specified IP options and then allow the packet to pass.

IP Options inspection can check for the following three IP options in a packet:

- End of Options List (EOOL) or IP Option 0—This option, which contains just a single zero byte, appears at the end of all options to mark the end of a list of options. This might not coincide with the end of the header according to the header length.

- No Operation (NOP) or IP Option 1—The Options field in the IP header can contain zero, one, or more options, which makes the total length of the field variable. However, the IP header must be a multiple of 32 bits. If the number of bits of all options is not a multiple of 32 bits, the NOP option is used as "internal padding" to align the options on a 32-bit boundary.

- Router Alert (RTRALT) or IP Option 20—This option notifies transit routers to inspect the contents of the packet even when the packet is not destined for that router. This inspection is valuable when implementing RSVP and similar protocols require relatively complex processing from the routers along the packets delivery path.

> **Note** IP Options inspection is included by default in the global inspection policy. Therefore, the ASA allows RSVP traffic that contains packets with the Router Alert option (option 20) when the ASA is in routed mode.

Dropping RSVP packets containing the Router Alert option can cause problems in VoIP implementations.

When you configure the ASA to clear the Router Alert option from IP headers, the IP header changes in the following ways:

- The Options field is padded so that the field ends on a 32 bit boundary.
- Internet header length (IHL) changes.
- The total length of the packet changes.
- The checksum is recomputed.

If an IP header contains additional options other than EOOL, NOP, or RTRALT, regardless of whether the ASA is configured to allow these options, the ASA will drop the packet.

Configuring IP Options Inspection

Use the Add Service Policy Rule Wizard - Rule Actions dialog box to configure IP Options inspection.

This wizard is available from the Configuration > Firewall > Service Policy Rules > Add > Add Service Policy Rule Wizard - Rule Actions dialog box.

Step 1   Open the Add Service Policy Rule Wizard by selecting **Configuration > Firewall > Service Policy Rules > Add**.

Perform the steps to complete the Service Policy, Traffic Classification Criteria, and Traffic Match - Destination Port pages of the wizard. See Adding a Service Policy Rule for Through Traffic, page 1-8.

The Add Service Policy Rule Wizard - Rule Actions dialog box opens.

Step 2   Check the **IP-Options** check box.

Step 3   Click **Configure**.

The Select IP Options Inspect Map dialog box opens.

Step 4   Perform one of the following:

- Click the **Use the default IP-Options inspection map** radio button to use the default IP Options map. The default map drops packets containing all the inspected IP options, namely End of Options List (EOOL), No Operation (NOP), and Router Alert (RTRALT).

- Click the **Select an IP-Options inspect map for fine control over inspection** radio button to select a defined application inspection map.

- Click Add to open the Add IP-Options Inspect Map dialog box and create a new inspection map.

Step 5   (Optional) If you clicked **Add** to create a new inspection map, define the following values for IP Options Inspection:

a.   Enter a name for the inspection map.

b.   Enter a description for the inspection map, up to 200 characters long.

c.   From the Parameters area, select which IP options you want to pass through the ASA or clear and then pass through the ASA:

   – Allow packets with the End of Options List (EOOL) option

   This option, which contains just a single zero byte, appears at the end of all options to mark the end of a list of options. This might not coincide with the end of the header according to the header length.

   – Allow packets with the No Operation (NOP) option

   The Options field in the IP header can contain zero, one, or more options, which makes the total length of the field variable. However, the IP header must be a multiple of 32 bits. If the number of bits of all options is not a multiple of 32 bits, the NOP option is used as "internal padding" to align the options on a 32-bit boundary.

   – Allow packets with the Router Alert (RTRALT) option

   This option notifies transit routers to inspect the contents of the packet even when the packet is not destined for that router. This inspection is valuable when implementing RSVP and similar protocols require relatively complex processing from the routers along the packets delivery path.

   – Clear the option value from the packets

   When an option is checked, the **Clear the option value from the packets** check box becomes available for that option. Select the **Clear the option value from the packets** check box to clear the option from the packet before allowing the packet through the ASA.

d.   Click **OK**.

Step 6   Click **OK**.

Step 7   Click **Finish**.

Select IP Options Inspect Map

The Select IP Options Inspect Map dialog box is accessible as follows:

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select IM Map

The Select IP-Options Inspect Map dialog box lets you select or create a new IP Options inspection map. Use this inspection map to control whether the ASA drops, passes, or clears IP packets containing the following IP options—End of Options List, No Operations, and Router Alert.

Fields

- Use the default IP-Options inspection map—Specifies to use the default IP Options map. The default map drops packets containing all the inspected IP options, namely End of Options List (EOOL), No Operation (NOP), and Router Alert (RTRALT).

- Select an IP-Options map for fine control over inspection**—**Lets you select a defined application inspection map or add a new one.

- Add—Opens the Add IP Options Inspect Map dialog box for the inspection.

IP Options Inspect Map

The IP Options Inspect Maps pane lets you view previously configured IP Options inspection maps. An IP Options inspection map lets you change the default configuration values used for IP Option inspection.

You can configure IP Options inspection to control which IP packets with specific IP options are allowed through the security appliance. Configuring this inspection instructs the security appliance to allow a packet to pass or to clear the specified IP options and then allow the packet to pass.

In particular, you can control whether the security appliance drops, clears, or passes packets containing the Router Alert (RTRALT) option. Dropping RSVP packets containing the Router Alert option can cause problems in VoIP implementations. Therefore, you can create IP Options inspection maps to pass packets containing the RTRALT option.

Fields

IP Options Inspect Maps—Table that lists the defined IP Options inspect maps.

Add—Configures a new IP Options inspect map.

Edit—Edits an existing IP Options inspect map. To edit an IP Options inspect map, choose the entry in the table and click Edit.

Delete—Deletes the inspect map selected in the IP Options Inspect Maps table.

Add/Edit IP Options Inspect Map

The Add/Edit IP Options Inspect Map lets you configure the settings for IP Options inspection maps.

Fields

- Name—When adding an IP Options inspection map, enter the name of the map. When editing a map, the name of the previously configured map is shown.

- Description—Enter the description of the IP Options inspection map, up to 200 characters in length.

- Parameters—Select which IP options you want to pass through the ASA or clear and then pass through the ASA:

  – Allow packets with the End of Options List (EOOL) option

  This option, which contains just a single zero byte, appears at the end of all options to mark the end of a list of options. This might not coincide with the end of the header according to the header length.

  – Allow packets with the No Operation (NOP) option

  The Options field in the IP header can contain zero, one, or more options, which makes the total length of the field variable. However, the IP header must be a multiple of 32 bits. If the number of bits of all options is not a multiple of 32 bits, the NOP option is used as "internal padding" to align the options on a 32-bit boundary.

  – Allow packets with the Router Alert (RTRALT) option

This option notifies transit routers to inspect the contents of the packet even when the packet is not destined for that router. This inspection is valuable when implementing RSVP and similar protocols require relatively complex processing from the routers along the packets delivery path.

  – Clear the option value from the packets

When an option is checked, the **Clear the option value from the packets** check box becomes available for that option. Select the **Clear the option value from the packets** check box to clear the option from the packet before allowing the packet through the ASA.

IPsec Pass Through Inspection

This section describes the IPsec Pass Through inspection engine. This section includes the following topics:

IPsec Pass Through Inspection Overview

Internet Protocol Security (IPsec) is a protocol suite for securing IP communications by authenticating and encrypting each IP packet of a data stream. IPsec also includes protocols for establishing mutual authentication between agents at the beginning of the session and negotiation of cryptographic keys to be used during the session. IPsec can be used to protect data flows between a pair of hosts (for example, computer users or servers), between a pair of security gateways (such as routers or firewalls), or between a security gateway and a host.

IPsec Pass Through application inspection provides convenient traversal of ESP (IP protocol 50) and AH (IP protocol 51) traffic associated with an IKE UDP port 500 connection. It avoids lengthy ACL configuration to permit ESP and AH traffic and also provides security using timeout and max connections.

Specify IPsec Pass Through inspection parameters to identify a specific map to use for defining the parameters for the inspection. Configure a policy map for Specify IPsec Pass Through inspection to access the parameters configuration, which lets you specify the restrictions for ESP or AH traffic. You can set the per client max connections and the idle timeout in parameters configuration.

NAT and non-NAT traffic is permitted. However, PAT is not supported.

Select IPsec-Pass-Thru Map

The Select IPsec-Pass-Thru Map dialog box is accessible as follows:

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select IPsec-Pass-Thru Map

The Select IPsec-Pass-Thru dialog box lets you select or create a new IPsec map. An IPsec map lets you change the configuration values used for IPsec application inspection. The Select IPsec Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default IPsec inspection map—Specifies to use the default IPsec map.

- Select an IPsec map for fine control over inspection—Lets you select a defined application inspection map or add a new one.

- Add—Opens the Add Policy Map dialog box for the inspection.

IPsec Pass Through Inspect Map

The IPsec Pass Through Inspect Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > IPsec Pass Through

The IPsec Pass Through pane lets you view previously configured IPsec Pass Through application inspection maps. An IPsec Pass Through map lets you change the default configuration values used for IPsec Pass Through application inspection. You can use an IPsec Pass Through map to permit certain flows without using an ACL.

Fields

- IPsec Pass Through Inspect Maps—Table that lists the defined IPsec Pass Through inspect maps.

- Add—Configures a new IPsec Pass Through inspect map. To edit an IPsec Pass Through inspect map, select the IPsec Pass Through entry in the IPsec Pass Through Inspect Maps table and click Customize.

- Delete—Deletes the inspect map selected in the IPsec Pass Through Inspect Maps table.

- Security Level—Select the security level (high or low).

    - Low—Default.

        Maximum ESP flows per client: Unlimited.

        ESP idle timeout: 00:10:00.

        Maximum AH flows per client: Unlimited.

        AH idle timeout: 00:10:00.

    - High

        Maximum ESP flows per client:10.

        ESP idle timeout: 00:00:30.

        Maximum AH flows per client: 10.

        AH idle timeout: 00:00:30.

    - Customize—Opens the Add/Edit IPsec Pass Thru Policy Map dialog box for additional settings.

    - Default Level—Sets the security level back to the default level of Low.

Add/Edit IPsec Pass Thru Policy Map (Security Level)

The Add/Edit IPsec Pass Thru Policy Map (Security Level) dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > IPsec Pass Through>
IPsec Pass Through Inspect Map > Basic View

The Add/Edit IPsec Pass Thru Policy Map pane lets you configure the security level and additional settings for IPsec Pass Thru application inspection maps.

Fields

- Name—When adding an IPsec Pass Thru map, enter the name of the IPsec Pass Thru map. When editing an IPsec Pass Thru map, the name of the previously configured IPsec Pass Thru map is shown.

- Security Level—Select the security level (high or low).

  - Low—Default.

    Maximum ESP flows per client: Unlimited.

    ESP idle timeout: 00:10:00.

    Maximum AH flows per client: Unlimited.

    AH idle timeout: 00:10:00.

  - High

    Maximum ESP flows per client:10.

    ESP idle timeout: 00:00:30.

    Maximum AH flows per client: 10.

    AH idle timeout: 00:00:30.

  - Default Level—Sets the security level back to the default level of Low.

- Details—Shows additional parameter settings to configure.

Add/Edit IPsec Pass Thru Policy Map (Details)

The Add/Edit IPsec Pass Thru Policy Map (Details) dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps >IPsec Pass Through > IPsec Pass Through Inspect Map > Advanced View

The Add/Edit IPsec Pass Thru Policy Map pane lets you configure the security level and additional settings for IPsec Pass Thru application inspection maps.

Fields

- Name—When adding an IPsec Pass Thru map, enter the name of the IPsec Pass Thru map. When editing an IPsec Pass Thru map, the name of the previously configured IPsec Pass Thru map is shown.

- Description—Enter the description of the IPsec Pass Through map, up to 200 characters in length.

- Security Level—Shows the security level settings to configure.

- Parameters—Configures ESP and AH parameter settings.

  - Limit ESP flows per client—Limits ESP flows per client.

    Maximum—Specify maximum limit.

  - Apply ESP idle timeout—Applies ESP idle timeout.

    Timeout—Specify timeout.

  - Limit AH flows per client—Limits AH flows per client.

Maximum—Specify maximum limit.

– Apply AH idle timeout—Applies AH idle timeout.

Timeout—Specify timeout.

IPv6 Inspection

Information about IPv6 Inspection

IPv6 inspection lets you selectively log or drop IPv6 traffic based on the extension header. In addition, IPv6 inspection can check conformance to RFC 2460 for type and order of extension headers in IPv6 packets.

Default Settings for IPv6 Inspection

If you enable IPv6 inspection and do not specify an inspection policy map, then the default IPv6 inspection policy map is used, and the following actions are taken:

- Allows only known IPv6 extension headers
- Enforces the order of IPv6 extension headers as defined in the RFC 2460 specification

If you create an inspection policy map, the above actions are taken by default unless you explicitly disable them.

(Optional) Configuring an IPv6 Inspection Policy Map

To identify extension headers to drop or log, and/or to disable packet verification, create an IPv6 inspection policy map to be used by the service policy.

Detailed Steps

Step 1   Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **IPv6**. The Configure IPv6 Maps pane appears.

Step 2   Click **Add**. The Add IPv6 Inspection Map dialog box appears.

Step 3   Enter a name and description for the inspection map.

By default, the Enforcement tab is selected and the following options are selected:

- Permit only known extension headers
- Enforce extension header order

When **Permit only known extension headers** is selected, the ASA verifies the IPv6 extension header.

When **Enforce extension header order** is selected, the order of IPv6 extension headers as defined in the RFC 2460 Specification is enforced.

When these options are specified and an error is detected, the ASA drops the packet and logs the action.

Step 4 To configure matching in the extension header, click the **Header Matches** tab.

Step 5 Click **Add** to add a match. The Add IPv6 Inspect dialog box appears.

    a. Select a criterion for the match.

    When you select any of the following criteria, you can configure to the ASA to drop or log when an IPv6 packet arrives matching the criterion:

      – Authentication (AH) header

      – Destination Options header

      – Encapsulating Security Payload (ESP) header

      – Fragment header

      – Hop-by-Hop Options header

      – Routing header—When Routing header is selected and an IPv6 routing extension header is detected, the ASA takes the specified action when the routing type is matched or a number when the specified routing type range is matched.

      – Header count—When Header count is selected and an IPv6 routing extension header is detected, the ASA takes the specified action when number of IPv6 extension headers in the packet is more than the specified value.

      – Routing header address count—When Routing header address count is selected, and an IPv6 routing extension header is detected, the ASA takes the specified action when the number of addresses in the type 0 routing header is more than the value you configure.

    b. Click **OK** to save the match criterion.

Step 6 Repeat Step 5 for each header you want to match.

Step 7 Click **OK** to save the IPv6 inspect map.

Configuring IPv6 Inspection

To enable IPv6 inspection, perform the following steps.

Detailed Steps

Step 1 Configure a service policy on the Configuration > Firewall > Service Policy Rules pane according to Chapter 1, "Service Policy."

You can configure IPv6 inspection as part of a new service policy rule, or you can edit an existing service policy.

Step 2 On the Rule Actions dialog box, click the **Protocol Inspections** tab.

Step 3 Check the **IPv6** check box.

Step 4 (Optional) To add an IPv6 inspection policy map that you configured in the (Optional) Configuring an IPv6 Inspection Policy Map, page 9-48:

    a. Click **Configure**.

The Select IPv6 Inspect Map dialog box appears.

b. Select the map name, and click **OK**.

Alternatively, you can click the **Add** button to add a new inspection policy map.

Step 5 Click **OK** or **Finish**.

NetBIOS Inspection

This section describes the IM inspection engine. This section includes the following topics:

NetBIOS Inspection Overview

NetBIOS inspection is enabled by default. The NetBios inspection engine translates IP addresses in the NetBios name service (NBNS) packets according to the ASA NAT configuration.

Select NETBIOS Map

The Select NETBIOS Map dialog box is accessible as follows:

**Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab >
Select NetBIOS Map**

The Select NETBIOS Map dialog box lets you select or create a new NetBIOS map. A NetBIOS map lets you change the configuration values used for NetBIOS application inspection. The Select NetBIOS Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default IM inspection map—Specifies to use the default NetBIOS map.
- Select a NetBIOS map for fine control over inspection**—**Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

NetBIOS Inspect Map

The NetBIOS Inspect Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > NetBIOS

The NetBIOS pane lets you view previously configured NetBIOS application inspection maps. A NetBIOS map lets you change the default configuration values used for NetBIOS application inspection.

NetBIOS application inspection performs NAT for the embedded IP address in the NetBIOS name service packets and NetBIOS datagram services packets. It also enforces protocol conformance, checking the various count and length fields for consistency.

Fields

- NetBIOS Inspect Maps—Table that lists the defined NetBIOS inspect maps.
- Add—Configures a new NetBIOS inspect map.
- Edit—Edits the selected NetBIOS entry in the NetBIOS Inspect Maps table.
- Delete—Deletes the inspect map selected in the NetBIOS Inspect Maps table.

Add/Edit NetBIOS Policy Map

The Add/Edit NetBIOS Policy Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > NetBIOS > NetBIOS Inspect Map > View

The Add/Edit NetBIOS Policy Map pane lets you configure the protocol violation settings for NetBIOS application inspection maps.

Fields

- Name—When adding a NetBIOS map, enter the name of the NetBIOS map. When editing an NetBIOS map, the name of the previously configured NetBIOS map is shown.
- Description—Enter the description of the NetBIOS map, up to 200 characters in length.
- Check for protocol violations—Checks for protocol violations and executes specified action.
    - Action—Drop packet or log.
    - Log—Enable or disable.

PPTP Inspection

PPTP is a protocol for tunneling PPP traffic. A PPTP session is composed of one TCP channel and usually two PPTP GRE tunnels. The TCP channel is the control channel used for negotiating and managing the PPTP GRE tunnels. The GRE tunnels carries PPP sessions between the two hosts.

When enabled, PPTP application inspection inspects PPTP protocol packets and dynamically creates the GRE connections and xlates necessary to permit PPTP traffic. Only Version 1, as defined in RFC 2637, is supported.

PAT is only performed for the modified version of GRE [RFC 2637] when negotiated over the PPTP TCP control channel. Port Address Translation is *not* performed for the unmodified version of GRE [RFC 1701, RFC 1702].

Specifically, the ASA inspects the PPTP version announcements and the outgoing call request/response sequence. Only PPTP Version 1, as defined in RFC 2637, is inspected. Further inspection on the TCP control channel is disabled if the version announced by either side is not Version 1. In addition, the outgoing-call request and reply sequence are tracked. Connections and xlates are dynamic allocated as necessary to permit subsequent secondary GRE data traffic.

The PPTP inspection engine must be enabled for PPTP traffic to be translated by PAT. Additionally, PAT is only performed for a modified version of GRE (RFC2637) and only if it is negotiated over the PPTP TCP control channel. PAT is not performed for the unmodified version of GRE (RFC 1701 and RFC 1702).

As described in RFC 2637, the PPTP protocol is mainly used for the tunneling of PPP sessions initiated from a modem bank PAC (PPTP Access Concentrator) to the headend PNS (PPTP Network Server). When used this way, the PAC is the remote client and the PNS is the server.

However, when used for VPN by Windows, the interaction is inverted. The PNS is a remote single-user PC that initiates connection to the head-end PAC to gain access to a central network.

SMTP and Extended SMTP Inspection

This section describes the IM inspection engine. This section includes the following topics:

SMTP and ESMTP Inspection Overview

ESMTP application inspection provides improved protection against SMTP-based attacks by restricting the types of SMTP commands that can pass through the ASA and by adding monitoring capabilities.

ESMTP is an enhancement to the SMTP protocol and is similar is most respects to SMTP. For convenience, the term SMTP is used in this document to refer to both SMTP and ESMTP. The application inspection process for extended SMTP is similar to SMTP application inspection and includes support for SMTP sessions. Most commands used in an extended SMTP session are the same as those used in an SMTP session but an ESMTP session is considerably faster and offers more options related to reliability and security, such as delivery status notification.

Extended SMTP application inspection adds support for these extended SMTP commands, including AUTH, EHLO, ETRN, HELP, SAML, SEND, SOML, STARTTLS, and VRFY. Along with the support for seven RFC 821 commands (DATA, HELO, MAIL, NOOP, QUIT, RCPT, RSET), the ASA supports a total of fifteen SMTP commands.

Other extended SMTP commands, such as ATRN, ONEX, VERB, CHUNKING, and private extensions and are not supported. Unsupported commands are translated into Xs, which are rejected by the internal server. This results in a message such as "500 Command unknown: 'XXX'." Incomplete commands are discarded.

The ESMTP inspection engine changes the characters in the server SMTP banner to asterisks except for the "2", "0", "0" characters. Carriage return (CR) and linefeed (LF) characters are ignored.

With SMTP inspection enabled, a Telnet session used for interactive SMTP may hang if the following rules are not observed: SMTP commands must be at least four characters in length; must be terminated with carriage return and line feed; and must wait for a response before issuing the next reply.

An SMTP server responds to client requests with numeric reply codes and optional human-readable strings. SMTP application inspection controls and reduces the commands that the user can use as well as the messages that the server returns. SMTP inspection performs three primary tasks:

- Restricts SMTP requests to seven basic SMTP commands and eight extended commands.
- Monitors the SMTP command-response sequence.
- Generates an audit trail—Audit record 108002 is generated when invalid character embedded in the mail address is replaced. For more information, see RFC 821.

SMTP inspection monitors the command and response sequence for the following anomalous signatures:

- Truncated commands.
- Incorrect command termination (not terminated with <CR><LR>).
- The MAIL and RCPT commands specify who are the sender and the receiver of the mail. Mail addresses are scanned for strange characters. The pipeline character (|) is deleted (changed to a blank space) and "<" ,">" are only allowed if they are used to define a mail address (">" must be preceded by "<").
- Unexpected transition by the SMTP server.
- For unknown commands, the ASA changes all the characters in the packet to X. In this case, the server generates an error code to the client. Because of the change in the packed, the TCP checksum has to be recalculated or adjusted.
- TCP stream editing.
- Command pipelining.

Select ESMTP Map

The Select ESMTP Map dialog box is accessible as follows:

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab >Select ESMTP Map

The Select ESMTP Map dialog box lets you select or create a new ESMTP map. An ESMTP map lets you change the configuration values used for ESMTP application inspection. The Select ESMTP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default ESMTP inspection map—Specifies to use the default ESMTP map.
- Select an ESMTP map for fine control over inspection**—**Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

ESMTP Inspect Map

The ESMTP Inspect Map dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > ESMTP

The ESMTP pane lets you view previously configured ESMTP application inspection maps. An ESMTP map lets you change the default configuration values used for ESMTP application inspection.

Since ESMTP traffic can be a main source of attack from spam, phising, malformed messages, buffer overflows, and buffer underflows, detailed packet inspection and control of ESMTP traffic are supported. Application security and protocol conformance enforce the sanity of the ESMTP message as well as detect several attacks, block senders and receivers, and block mail relay.

Fields

- ESMTP Inspect Maps—Table that lists the defined ESMTP inspect maps.

- Add—Configures a new ESMTP inspect map. To edit an ESMTP inspect map, choose the ESMTP entry in the ESMTP Inspect Maps table and click **Customize**.

- Delete—Deletes the inspect map selected in the ESMTP Inspect Maps table.

- Security Level—Select the security level (high, medium, or low).

    - Low—Default.

        Log if command line length is greater than 512

        Log if command recipient count is greater than 100

        Log if body line length is greater than 1000

        Log if sender address length is greater than 320

        Log if MIME file name length is greater than 255

    - Medium

        Obfuscate Server Banner

        Drop Connections if command line length is greater than 512

        Drop Connections if command recipient count is greater than 100

        Drop Connections if body line length is greater than 1000

        Drop Connections if sender address length is greater than 320

        Drop Connections if MIME file name length is greater than 255

    - High

        Obfuscate Server Banner

        Drop Connections if command line length is greater than 512

        Drop Connections if command recipient count is greater than 100

        Drop Connections if body line length is greater than 1000

        Drop Connections and log if sender address length is greater than 320

        Drop Connections and log if MIME file name length is greater than 255

    - MIME File Type Filtering—Opens the MIME Type Filtering dialog box to configure MIME file type filters.

    - Customize—Opens the Add/Edit ESMTP Policy Map dialog box for additional settings.

    - Default Level—Sets the security level back to the default level of Low.

MIME File Type Filtering

The MIME File Type Filtering dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > ESMTP > MIME File Type Filtering

The MIME File Type Filtering dialog box lets you configure the settings for a MIME file type filter.

Fields

- Match Type—Shows the match type, which can be a positive or negative match.
- Criterion—Shows the criterion of the inspection.
- Value—Shows the value to match in the inspection.
- Action—Shows the action if the match condition is met.
- Log—Shows the log state.
- Add—Opens the Add MIME File Type Filter dialog box to add a MIME file type filter.
- Edit—Opens the Edit MIME File Type Filter dialog box to edit a MIME file type filter.
- Delete—Deletes a MIME file type filter.
- Move Up—Moves an entry up in the list.
- Move Down—Moves an entry down in the list.

Add/Edit ESMTP Policy Map (Security Level)

The Add/Edit ESMTP Policy Map (Security Level) dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > ESMTP > ESMTP Inspect Map > Basic View

The Add/Edit ESMTP Policy Map pane lets you configure the security level and additional settings for ESMTP application inspection maps.

Fields

- Name—When adding an ESMTP map, enter the name of the ESMTP map. When editing an ESMTP map, the name of the previously configured ESMTPS map is shown.
- Description—Enter the description of the ESMTP map, up to 200 characters in length.
- Security Level—Select the security level (high, medium, or low).
    - Low—Default.

        Log if command line length is greater than 512

        Log if command recipient count is greater than 100

        Log if body line length is greater than 1000

        Log if sender address length is greater than 320

        Log if MIME file name length is greater than 255

    - Medium

        Obfuscate Server Banner

        Drop Connections if command line length is greater than 512

        Drop Connections if command recipient count is greater than 100

        Drop Connections if body line length is greater than 1000

        Drop Connections if sender address length is greater than 320

        Drop Connections if MIME file name length is greater than 255

    - High

Obfuscate Server Banner

Drop Connections if command line length is greater than 512

Drop Connections if command recipient count is greater than 100

Drop Connections if body line length is greater than 1000

Drop Connections and log if sender address length is greater than 320

Drop Connections and log if MIME file name length is greater than 255

- – MIME File Type Filtering—Opens the MIME Type Filtering dialog box to configure MIME file type filters.

- – Default Level—Sets the security level back to the default level of Low.

- • Details—Shows the Parameters and Inspections tabs to configure additional settings.

Add/Edit ESMTP Policy Map (Details)

The Add/Edit ESMTP Policy Map (Details) dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > ESMTP > ESMTP Inspect Map > Advanced View

The Add/Edit ESMTP Policy Map pane lets you configure the security level and additional settings for ESMTP application inspection maps.

Fields

- • Name—When adding an ESMTP map, enter the name of the ESMTP map. When editing an ESMTP map, the name of the previously configured ESMTP map is shown.

- • Description—Enter the description of the ESMTP map, up to 200 characters in length.

- • Security Level—Shows the security level and mime file type filtering settings to configure.

- • Parameters—Tab that lets you configure the parameters for the ESMTP inspect map.

  - – Mask server banner—Enforces banner obfuscation.

  - – Configure Mail Relay—Enables ESMTP mail relay.

    Domain Name—Specifies a local domain.

    Action—Drop connection or log.

    Log—Enable or disable.

- • Inspections—Tab that shows you the ESMTP inspection configuration and lets you add or edit.

  - – Match Type—Shows the match type, which can be a positive or negative match.

  - – Criterion—Shows the criterion of the ESMTP inspection.

  - – Value—Shows the value to match in the ESMTP inspection.

  - – Action—Shows the action if the match condition is met.

  - – Log—Shows the log state.

  - – Add—Opens the Add ESMTP Inspect dialog box to add an ESMTP inspection.

  - – Edit—Opens the Edit ESMTP Inspect dialog box to edit an ESMTP inspection.

  - – Delete—Deletes an ESMTP inspection.

  - – Move Up—Moves an inspection up in the list.

- Move Down—Moves an inspection down in the list.

Add/Edit ESMTP Inspect

The Add/Edit ESMTP Inspect dialog box is accessible as follows:

Configuration > Global Objects > Inspect Maps > ESMTP > ESMTP Inspect Map > Advanced View > Add/Edit ESMTP Inspect

The Add/Edit ESMTP Inspect dialog box lets you define the match criterion and value for the ESMTP inspect map.

Fields

- Match Type—Specifies whether traffic should match or not match the values.

  For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map.

- Criterion—Specifies which criterion of ESMTP traffic to match.

  - Body Length—Match body length at specified length in bytes.
  - Body Line Length—Match body line length matching at specified length in bytes.
  - Commands—Match commands exchanged in the ESMTP protocol.
  - Command Recipient Count—Match command recipient count greater than number specified.
  - Command Line Length—Match command line length greater than length specified in bytes.
  - EHLO Reply Parameters—Match an ESMTP ehlo reply parameter.
  - Header Length—Match header length at length specified in bytes.
  - Header To Fields Count—Match header To fields count greater than number specified.
  - Invalid Recipients Count—Match invalid recipients count greater than number specified.
  - MIME File Type—Match MIME file type.
  - MIME Filename Length—Match MIME filename.
  - MIME Encoding—Match MIME encoding.
  - Sender Address—Match sender email address.
  - Sender Address Length—Match sender email address length.

- Body Length Criterion Values—Specifies the value details for body length match.

  - Greater Than Length—Body length in bytes.
  - Action—Reset, drop connection, log.
  - Log—Enable or disable.

- Body Line Length Criterion Values—Specifies the value details for body line length match.

  - Greater Than Length—Body line length in bytes.
  - Action—Reset, drop connection, log.
  - Log—Enable or disable.

- Commands Criterion Values—Specifies the value details for command match.

  - Available Commands Table:

AUTH

DATA

EHLO

ETRN

HELO

HELP

MAIL

NOOP

QUIT

RCPT

RSET

SAML

SOML

VRFY

- – Add—Adds the selected command from the Available Commands table to the Selected Commands table.

- – Remove—Removes the selected command from the Selected Commands table.

- – Primary Action—Mask, Reset, Drop Connection, None, Limit Rate (pps).

- – Log—Enable or disable.

- – Rate Limit—Do not limit rate, Limit Rate (pps).

- • Command Recipient Count Criterion Values—Specifies the value details for command recipient count match.

- – Greater Than Count—Specify command recipient count.

- – Action—Reset, drop connection, log.

- – Log—Enable or disable.

- • Command Line Length Criterion Values—Specifies the value details for command line length.

- – Greater Than Length—Command line length in bytes.

- – Action—Reset, drop connection, log.

- – Log—Enable or disable.

- • EHLO Reply Parameters Criterion Values—Specifies the value details for EHLO reply parameters match.

- – Available Parameters Table:

    8bitmime

    auth

    binarymime

    checkpoint

    dsn

    ecode

etrn

others

pipelining

size

vrfy

- **–** Add—Adds the selected parameter from the Available Parameters table to the Selected Parameters table.

- **–** Remove—Removes the selected command from the Selected Commands table.

- **–** Action—Reset, Drop Connection, Mask, Log.

- **–** Log—Enable or disable.

- Header Length Criterion Values—Specifies the value details for header length match.

  - **–** Greater Than Length—Header length in bytes.

  - **–** Action—Reset, Drop Connection, Mask, Log.

  - **–** Log—Enable or disable.

- Header To Fields Count Criterion Values—Specifies the value details for header To fields count match.

  - **–** Greater Than Count—Specify command recipient count.

  - **–** Action—Reset, drop connection, log.

  - **–** Log—Enable or disable.

- Invalid Recipients Count Criterion Values—Specifies the value details for invalid recipients count match.

  - **–** Greater Than Count—Specify command recipient count.

  - **–** Action—Reset, drop connection, log.

  - **–** Log—Enable or disable.

- MIME File Type Criterion Values—Specifies the value details for MIME file type match.

  - **–** Regular Expression—Lists the defined regular expressions to match.

  - **–** Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  - **–** Regular Expression Class—Lists the defined regular expression classes to match.

  - **–** Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

  - **–** Action—Reset, drop connection, log.

  - **–** Log—Enable or disable.

- MIME Filename Length Criterion Values—Specifies the value details for MIME filename length match.

  - **–** Greater Than Length—MIME filename length in bytes.

  - **–** Action—Reset, Drop Connection, Log.

  - **–** Log—Enable or disable.

- MIME Encoding Criterion Values—Specifies the value details for MIME encoding match.

- – Available Encodings table

  7bit

  8bit

  base64

  binary

  others

  quoted-printable

  – Add—Adds the selected parameter from the Available Encodings table to the Selected Encodings table.

  – Remove—Removes the selected command from the Selected Commands table.

  – Action—Reset, Drop Connection, Log.

  – Log—Enable or disable.

- • Sender Address Criterion Values—Specifies the value details for sender address match.

  – Regular Expression—Lists the defined regular expressions to match.

  – Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  – Regular Expression Class—Lists the defined regular expression classes to match.

  – Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

  – Action—Reset, Drop Connection, Log.

  – Log—Enable or disable.

- • Sender Address Length Criterion Values—Specifies the value details for sender address length match.

  – Greater Than Length—Sender address length in bytes.

  – Action—Reset, Drop Connection, Log.

  – Log—Enable or disable.

TFTP Inspection

TFTP inspection is enabled by default.

TFTP, described in RFC 1350, is a simple protocol to read and write files between a TFTP server and client.

The ASA inspects TFTP traffic and dynamically creates connections and translations, if necessary, to permit file transfer between a TFTP client and server. Specifically, the inspection engine inspects TFTP read request (RRQ), write request (WRQ), and error notification (ERROR).

A dynamic secondary channel and a PAT translation, if necessary, are allocated on a reception of a valid read (RRQ) or write (WRQ) request. This secondary channel is subsequently used by TFTP for file transfer or error notification.

Only the TFTP server can initiate traffic over the secondary channel, and at most one incomplete secondary channel can exist between the TFTP client and server. An error notification from the server closes the secondary channel.

TFTP inspection must be enabled if static PAT is used to redirect TFTP traffic.

Inspection for Voice and Video Protocols

This chapter describes how to configure application layer protocol inspection. Inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the ASA to do a deep packet inspection instead of passing the packet through the fast path. As a result, inspection engines can affect overall throughput.

Several common inspection engines are enabled on the ASA by default, but you might need to enable others depending on your network.

This chapter includes the following sections:

CTIQBE Inspection

This section describes CTIQBE application inspection. This section includes the following topics:

CTIQBE Inspection Overview

CTIQBE protocol inspection supports NAT, PAT, and bidirectional NAT. This enables Cisco IP SoftPhone and other Cisco TAPI/JTAPI applications to work successfully with Cisco CallManager for call setup across the ASA.

TAPI and JTAPI are used by many Cisco VoIP applications. CTIQBE is used by Cisco TSP to communicate with Cisco CallManager.

Limitations and Restrictions

The following summarizes limitations that apply when using CTIQBE application inspection:

- CTIQBE application inspection does not support configurations with the **alias** command.

- Stateful failover of CTIQBE calls is not supported.

- Debugging CTIQBE inspection may delay message transmission, which may have a performance impact in a real-time environment. When you enable this debugging or logging and Cisco IP SoftPhone seems unable to complete call setup through the ASA, increase the timeout values in the Cisco TSP settings on the system running Cisco IP SoftPhone.

The following summarizes special considerations when using CTIQBE application inspection in specific scenarios:

- If two Cisco IP SoftPhones are registered with different Cisco CallManagers, which are connected to different interfaces of the ASA, calls between these two phones fails.

- When Cisco CallManager is located on the higher security interface compared to Cisco IP SoftPhones, if NAT or outside NAT is required for the Cisco CallManager IP address, the mapping must be static as Cisco IP SoftPhone requires the Cisco CallManager IP address to be specified explicitly in its Cisco TSP configuration on the PC.

- When using PAT or Outside PAT, if the Cisco CallManager IP address is to be translated, its TCP port 2748 must be statically mapped to the same port of the PAT (interface) address for Cisco IP SoftPhone registrations to succeed. The CTIQBE listening port (TCP 2748) is fixed and is not user-configurable on Cisco CallManager, Cisco IP SoftPhone, or Cisco TSP.

H.323 Inspection

This section describes the H.323 application inspection. This section includes the following topics:

H.323 Inspection Overview

H.323 inspection provides support for H.323 compliant applications such as Cisco CallManager and VocalTec Gatekeeper. H.323 is a suite of protocols defined by the International Telecommunication Union for multimedia conferences over LANs. The ASA supports H.323 through Version 6, including H.323 v3 feature Multiple Calls on One Call Signaling Channel.

With H.323 inspection enabled, the ASA supports multiple calls on the same call signaling channel, a feature introduced with H.323 Version 3. This feature reduces call setup time and reduces the use of ports on the ASA.

The two major functions of H.323 inspection are as follows:

- NAT the necessary embedded IPv4 addresses in the H.225 and H.245 messages. Because H.323 messages are encoded in PER encoding format, the ASA uses an ASN.1 decoder to decode the H.323 messages.

- Dynamically allocate the negotiated H.245 and RTP/RTCP connections.

How H.323 Works

The H.323 collection of protocols collectively may use up to two TCP connection and four to eight UDP connections. FastConnect uses only one TCP connection, and RAS uses a single UDP connection for registration, admissions, and status.

An H.323 client can initially establish a TCP connection to an H.323 server using TCP port 1720 to request Q.931 call setup. As part of the call setup process, the H.323 terminal supplies a port number to the client to use for an H.245 TCP connection. In environments where H.323 gatekeeper is in use, the initial packet is transmitted using UDP.

H.323 inspection monitors the Q.931 TCP connection to determine the H.245 port number. If the H.323 terminals are not using FastConnect, the ASA dynamically allocates the H.245 connection based on the inspection of the H.225 messages.

Note   The H.225 connection can also be dynamically allocated when using RAS.

Within each H.245 message, the H.323 endpoints exchange port numbers that are used for subsequent UDP data streams. H.323 inspection inspects the H.245 messages to identify these ports and dynamically creates connections for the media exchange. RTP uses the negotiated port number, while RTCP uses the next higher port number.

The H.323 control channel handles H.225 and H.245 and H.323 RAS. H.323 inspection uses the following ports.

- 1718—Gate Keeper Discovery UDP port
- 1719—RAS UDP port
- 1720—TCP Control Port

You must permit traffic for the well-known H.323 port 1719 for RAS signaling. Additionally, you must permit traffic for the well-known H.323 port 1720 for the H.225 call signaling; however, the H.245 signaling ports are negotiated between the endpoints in the H.225 signaling. When an H.323 gatekeeper is used, the ASA opens an H.225 connection based on inspection of the ACF and RCF nmessages.

After inspecting the H.225 messages, the ASA opens the H.245 channel and then inspects traffic sent over the H.245 channel as well. All H.245 messages passing through the ASA undergo H.245 application inspection, which translates embedded IP addresses and opens the media channels negotiated in H.245 messages.

The H.323 ITU standard requires that a TPKT header, defining the length of the message, precede the H.225 and H.245, before being passed on to the reliable connection. Because the TPKT header does not necessarily need to be sent in the same TCP packet as H.225 and H.245 messages, the ASA must remember the TPKT length to process and decode the messages properly. For each connection, the ASA keeps a record that contains the TPKT length for the next expected message.

If the ASA needs to perform NAT on IP addresses in messages, it changes the checksum, the UUIE length, and the TPKT, if it is included in the TCP packet with the H.225 message. If the TPKT is sent in a separate TCP packet, the ASA proxy ACKs that TPKT and appends a new TPKT to the H.245 message with the new length.

Note The ASA does not support TCP options in the Proxy ACK for the TPKT.

Each UDP connection with a packet going through H.323 inspection is marked as an H.323 connection and times out with the H.323 timeout as configured in the Configuration > Firewall > Advanced > Global Timeouts pane.

Note You can enable call setup between H.323 endpoints when the Gatekeeper is inside the network. The ASA includes options to open pinholes for calls based on the RegistrationRequest/RegistrationConfirm (RRQ/RCF) messages. Because these RRQ/RCF messages are sent to and from the Gatekeeper, the calling endpoint's IP address is unknown and the ASA opens a pinhole through source IP address/port 0/0. By default, this option is disabled.

H.239 Support in H.245 Messages

The ASA sits between two H.323 endpoints. When the two H.323 endpoints set up a telepresentation session so that the endpoints can send and receive a data presentation, such as spreadsheet data, the ASA ensure successful H.239 negotiation between the endpoints.

H.239 is a standar that provides the ability for H.300 series endpoints to open an additional video channel in a single call. In a call, an endpoint (such as a video phone), sends a channel for video and a channel for data presentation. The H.239 negotiation occurs on the H.245 channel.

The ASA opens pinholes for the additional media channel and the media control channel. The endpoints use open logical channel message (OLC) to signal a new channel creation. The message extension is part of H.245 version 13.

The decoding and encoding of of the telepresentation session is enabled by default. H.239 encoding and decoding is preformed by ASN.1 coder.

Limitations and Restrictions

The following are some of the known issues and limitations when using H.323 application inspection:

• Only static NAT is fully supported. Static PAT may not properly translate IP addresses embedded in optional fields within H.323 messages. If you experience this kind of problem, do not use static PAT with H.323.

- Not supported with dynamic NAT or PAT.

- Not supported with extended PAT.

- Not supported with NAT between same-security-level interfaces.

- Not supported with outside NAT.

- Not supported with NAT64.

- When a NetMeeting client registers with an H.323 gatekeeper and tries to call an H.323 gateway that is also registered with the H.323 gatekeeper, the connection is established but no voice is heard in either direction. This problem is unrelated to the ASA.

- If you configure a network static address where the network static address is the same as a third-party netmask and address, then any outbound H.323 connection fails.

Select H.323 Map

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select H.323 Map

The Select H.323 Map dialog box lets you select or create a new H.323 map. An H.323 map lets you change the configuration values used for H.323 application inspection. The Select H.323 Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default H.323 inspection map—Specifies to use the default H.323 map.

- Select an H.323 map for fine control over inspection**—**Lets you select a defined application inspection map or add a new one.

- Add—Opens the Add Policy Map dialog box for the inspection.

H.323 Class Map

Configuration > Global Objects > Class Maps > H.323

The H.323 Class Map pane lets you configure H.323 class maps for H.323 inspection.

An inspection class map matches application traffic with criteria specific to the application. You then identify the class map in the inspect map and enable actions. The difference between creating a class map and defining the traffic match directly in the inspect map is that you can create more complex match criteria and you can reuse class maps. The applications that support inspection class maps are DNS, FTP, H.323, HTTP, IM, and SIP.

Fields

- Name—Shows the H.323 class map name.

- Match Conditions—Shows the type, match criterion, and value in the class map.

  - Match Type—Shows the match type, which can be a positive or negative match.

  - Criterion—Shows the criterion of the H.323 class map.

  - Value—Shows the value to match in the H.323 class map.

- Description—Shows the description of the class map.

- Add—Adds an H.323 class map.

- Edit—Edits an H.323 class map.
- Delete—Deletes an H.323 class map.

Add/Edit H.323 Traffic Class Map

Configuration > Global Objects > Class Maps > H.323 > Add/Edit H.323 Traffic Class Map

The Add/Edit H.323 Traffic Class Map dialog box lets you define a H.323 class map.

Fields

- Name—Enter the name of the H.323 class map, up to 40 characters in length.
- Description—Enter the description of the H.323 class map.
- Add—Adds an H.323 class map.
- Edit—Edits an H.323 class map.
- Delete—Deletes an H.323 class map.

Add/Edit H.323 Match Criterion

Configuration > Global Objects > Class Maps > H.323 > Add/Edit H.323 Traffic Class Map > Add/Edit H.323 Match Criterion

The Add/Edit H.323 Match Criterion dialog box lets you define the match criterion and value for the H.323 class map.

Fields

- Match Type—Specifies whether the class map includes traffic that matches the criterion, or traffic that does not match the criterion.

  For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map.

- Criterion—Specifies which criterion of H.323 traffic to match.
  - Called Party—Match the called party.
  - Calling Party—Match the calling party.
  - Media Type—Match the media type.
- Called Party Criterion Values—Specifies to match on the H.323 called party.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Calling Party Criterion Values—Specifies to match on the H.323 calling party.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

- Regular Expression Class—Lists the defined regular expression classes to match.

- Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Media Type Criterion Values—Specifies which media type to match.

  - Audio—Match audio type.

  - Video—Match video type.

  - Data—Match data type.

H.323 Inspect Map

Configuration > Global Objects > Inspect Maps > H.323

The H.323 pane lets you view previously configured H.323 application inspection maps. An H.323 map lets you change the default configuration values used for H.323 application inspection.

H.323 inspection supports RAS, H.225, and H.245, and its functionality translates all embedded IP addresses and ports. It performs state tracking and filtering and can do a cascade of inspect function activation. H.323 inspection supports phone number filtering, dynamic T.120 control, H.245 tunneling control, HSI groups, protocol state tracking, H.323 call duration enforcement, and audio/video control.

Fields

- H.323 Inspect Maps—Table that lists the defined H.323 inspect maps.

- Add—Configures a new H.323 inspect map. To edit an H.323 inspect map, choose the H.323 entry in the H.323 Inspect Maps table and click **Customize**.

- Delete—Deletes the inspect map selected in the H.323 Inspect Maps table.

- Security Level—Select the security level (low, medium, or high).

  - Low—Default.

    State Checking h225 Disabled

    State Checking ras Disabled

    Call Party Number Disabled

    Call duration Limit Disabled

    RTP conformance not enforced

  - Medium

    State Checking h225 Enabled

    State Checking ras Enabled

    Call Party Number Disabled

    Call duration Limit Disabled

    RTP conformance enforced

    Limit payload to audio or video, based on the signaling exchange: no

  - High

    State Checking h225 Enabled

    State Checking ras Enabled

Call Party Number Enabled

Call duration Limit 1:00:00

RTP conformance enforced

Limit payload to audio or video, based on the signaling exchange: yes

- Phone Number Filtering—Opens the Phone Number Filtering dialog box to configure phone number filters.

- Customize—Opens the Add/Edit H.323 Policy Map dialog box for additional settings.

- Default Level—Sets the security level back to the default level of Medium.

Phone Number Filtering

Configuration > Global Objects > Inspect Maps > H323 > Phone Number Filtering

The Phone Number Filtering dialog box lets you configure the settings for a phone number filter.

Fields

- Match Type—Shows the match type, which can be a positive or negative match.

- Criterion—Shows the criterion of the inspection.

- Value—Shows the value to match in the inspection.

- Action—Shows the action if the match condition is met.

- Log—Shows the log state.

- Add—Opens the Add Phone Number Filter dialog box to add a phone number filter.

- Edit—Opens the Edit Phone Number Filter dialog box to edit a phone number filter.

- Delete—Deletes a phone number filter.

- Move Up—Moves an entry up in the list.

- Move Down—Moves an entry down in the list.

Add/Edit H.323 Policy Map (Security Level)

Configuration > Global Objects > Inspect Maps > H323 > H323 Inspect Map > Basic View

The Add/Edit H.323 Policy Map pane lets you configure the security level and additional settings for H.323 application inspection maps.

Fields

- Name—When adding an H.323 map, enter the name of the H.323 map. When editing an H.323 map, the name of the previously configured H.323 map is shown.

- Description—Enter the description of the H323 map, up to 200 characters in length.

- Security Level—Select the security level (low, medium, or high).

  - Low—Default.

    State Checking h225 Disabled

    State Checking ras Disabled

Call Party Number Disabled

Call duration Limit Disabled

RTP conformance not enforced

– Medium

State Checking h225 Enabled

State Checking ras Enabled

Call Party Number Disabled

Call duration Limit Disabled

RTP conformance enforced

Limit payload to audio or video, based on the signaling exchange: no

– High

State Checking h225 Enabled

State Checking ras Enabled

Call Party Number Enabled

Call duration Limit 1:00:00

RTP conformance enforced

Limit payload to audio or video, based on the signaling exchange: yes

– Phone Number Filtering—Opens the Phone Number Filtering dialog box which lets you configure the settings for a phone number filter.

– Default Level—Sets the security level back to the default.

- Details—Shows the State Checking, Call Attributes, Tunneling and Protocol Conformance, HSI Group Parameters, and Inspections tabs to configure additional settings.

Add/Edit H.323 Policy Map (Details)

Configuration > Global Objects > Inspect Maps > H323 > H323 Inspect Map > Advanced View

The Add/Edit H.323 Policy Map pane lets you configure the security level and additional settings for H.323 application inspection maps.

Fields

- Name—When adding an H.323 map, enter the name of the H.323 map. When editing an H.323 map, the name of the previously configured H.323 map is shown.

- Description—Enter the description of the H.323 map, up to 200 characters in length.

- Security Level—Shows the security level and phone number filtering settings to configure.

- State Checking—Tab that lets you configure state checking parameters for the H.323 inspect map.

  – Check state transition of H.225 messages—Enforces H.323 state checking on H.225 messages.

  – Check state transition of RAS messages—Enforces H.323 state checking on RAS messages.

  – Check RFC messages and open pinholes for call signal addresses in RFQ messages

> **Note** You can enable call setup between H.323 endpoints when the Gatekeeper is inside the network. The ASA includes options to open pinholes for calls based on the RegistrationRequest/RegistrationConfirm (RRQ/RCF) messages. Because these RRQ/RCF messages are sent to and from the Gatekeeper, the calling endpoint's IP address is unknown and the ASA opens a pinhole through source IP address/port 0/0. By default, this option is disabled. You can enable this option by setting the option in the H.323 Inspect Map.

- Call Attributes—Tab that lets you configure call attributes parameters for the H.323 inspect map.
  - Enforce call duration limit—Enforces the absolute limit on a call.

    Call Duration Limit—Time limit for the call (hh:mm:ss).
  - Enforce presence of calling and called party numbers—Enforces sending call party numbers during call setup.

- Tunneling and Protocol Conformance—Tab that lets you configure tunneling and protocol conformance parameters for the H.323 inspect map.
  - Check for H.245 tunneling—Allows H.245 tunneling.

    Action—Drop connection or log.
  - Check RTP packets for protocol conformance—Checks RTP/RTCP packets on the pinholes for protocol conformance.

    Limit payload to audio or video, based on the signaling exchange—Enforces the payload type to be audio or video based on the signaling exchange.

- HSI Group Parameters—Tab that lets you configure an HSI group.
  - HSI Group ID—Shows the HSI Group ID.
  - IP Address—Shows the HSI Group IP address.
  - Endpoints—Shows the HSI Group endpoints.
  - Add—Opens the Add HSI Group dialog box to add an HSI group.
  - Edit—Opens the Edit HSI Group dialog box to edit an HSI group.
  - Delete—Deletes an HSI group.

- Inspections—Tab that shows you the H.323 inspection configuration and lets you add or edit.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the H.323 inspection.
  - Value—Shows the value to match in the H.323 inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add H.323 Inspect dialog box to add an H.323 inspection.
  - Edit—Opens the Edit H.323 Inspect dialog box to edit an H.323 inspection.
  - Delete—Deletes an H.323 inspection.
  - Move Up—Moves an inspection up in the list.
  - Move Down—Moves an inspection down in the list.

Add/Edit HSI Group

Configuration > Global Objects > Inspect Maps > H323 > H323 Inspect Map > Advanced View > Add/Edit HSI Group

The Add/Edit HSI Group dialog box lets you configure HSI Groups.

Fields

- Group ID—Enter the HSI group ID.
- IP Address—Enter the HSI IP address.
- Endpoints—Lets you configure the IP address and interface of the endpoints.
    - IP Address—Enter an endpoint IP address.
    - Interface—Specifies an endpoint interface.
- Add—Adds the HSI group defined.
- Delete—Deletes the selected HSI group.

Add/Edit H.323 Map

Configuration > Global Objects > Inspect Maps > H232 > H323 Inspect Map > Advanced View > Add/Edit H323 Inspect

The Add/Edit H.323 Inspect dialog box lets you define the match criterion and value for the H.323 inspect map.

Fields

- Single Match—Specifies that the H.323 inspect has only one match statement.
- Match Type—Specifies whether traffic should match or not match the values.

  For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map.

- Criterion—Specifies which criterion of H.323 traffic to match.
    - Called Party—Match the called party.
    - Calling Party—Match the calling party.
    - Media Type—Match the media type.
- Called Party Criterion Values—Specifies to match on the H.323 called party.
    - Regular Expression—Lists the defined regular expressions to match.
    - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
    - Regular Expression Class—Lists the defined regular expression classes to match.
    - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Calling Party Criterion Values—Specifies to match on the H.323 calling party.
    - Regular Expression—Lists the defined regular expressions to match.
    - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

- Regular Expression Class—Lists the defined regular expression classes to match.
- Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Media Type Criterion Values—Specifies which media type to match.
  - Audio—Match audio type.
  - Video—Match video type.
  - Data—Match data type.

- Multiple Matches—Specifies multiple matches for the H.323 inspection.
  - H323 Traffic Class—Specifies the H.323 traffic class match.
  - Manage—Opens the Manage H323 Class Maps dialog box to add, edit, or delete H.323 Class Maps.

- Action—Drop packet, drop connection, or reset.

MGCP Inspection

This section describes MGCP application inspection. This section includes the following topics:

MGCP Inspection Overview

MGCP is a master/slave protocol used to control media gateways from external call control elements called media gateway controllers or call agents. A media gateway is typically a network element that provides conversion between the audio signals carried on telephone circuits and data packets carried over the Internet or over other packet networks. Using NAT and PAT with MGCP lets you support a large number of devices on an internal network with a limited set of external (global) addresses. Examples of media gateways are:

- Trunking gateways, that interface between the telephone network and a Voice over IP network. Such gateways typically manage a large number of digital circuits.
- Residential gateways, that provide a traditional analog (RJ11) interface to a Voice over IP network. Examples of residential gateways include cable modem/cable set-top boxes, xDSL devices, broad-band wireless devices.
- Business gateways, that provide a traditional digital PBX interface or an integrated soft PBX interface to a Voice over IP network.

Note To avoid policy failure when upgrading from ASA version 7.1, all layer 7 and layer 3 policies must have distinct names. For instance, a previously configured policy map with the same name as a previously configured MGCP map must be changed before the upgrade.

MGCP messages are transmitted over UDP. A response is sent back to the source address (IP address and UDP port number) of the command, but the response may not arrive from the same address as the command was sent to. This can happen when multiple call agents are being used in a failover configuration and the call agent that received the command has passed control to a backup call agent, which then sends the response. Figure 10-1 illustrates how NAT can be used with MGCP.

Figure 10-1   Using NAT with MGCP
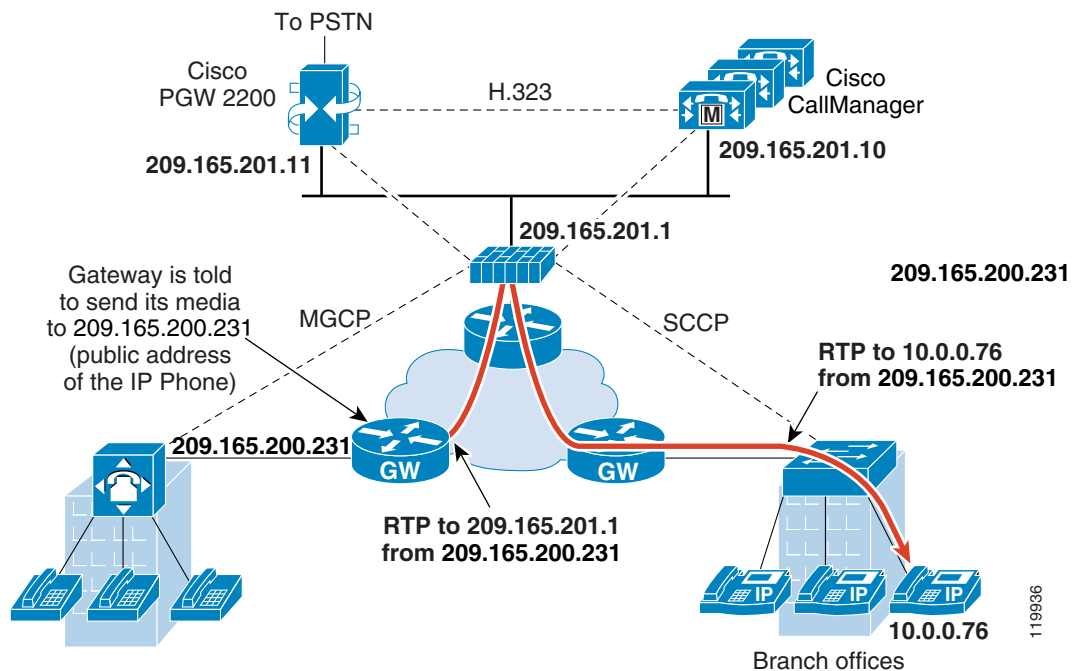
MGCP endpoints are physical or virtual sources and destinations for data. Media gateways contain endpoints on which the call agent can create, modify and delete connections to establish and control media sessions with other multimedia endpoints. Also, the call agent can instruct the endpoints to detect certain events and generate signals. The endpoints automatically communicate changes in service state to the call agent.

MGCP transactions are composed of a command and a mandatory response. There are eight types of commands:

- CreateConnection
- ModifyConnection
- DeleteConnection
- NotificationRequest
- Notify
- AuditEndpoint
- AuditConnection

- RestartInProgress

The first four commands are sent by the call agent to the gateway. The Notify command is sent by the gateway to the call agent. The gateway may also send a DeleteConnection. The registration of the MGCP gateway with the call agent is achieved by the RestartInProgress command. The AuditEndpoint and the AuditConnection commands are sent by the call agent to the gateway.

All commands are composed of a Command header, optionally followed by a session description. All responses are composed of a Response header, optionally followed by a session description.

- The port on which the gateway receives commands from the call agent. Gateways usually listen to UDP port 2427.

- The port on which the call agent receives commands from the gateway. Call agents usually listen to UDP port 2727.

> **Note** MGCP inspection does not support the use of different IP addresses for MGCP signaling and RTP data. A common and recommended practice is to send RTP data from a resilient IP address, such as a loopback or virtual IP address; however, the ASA requires the RTP data to come from the same address as MGCP signalling.

Select MGCP Map

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select MGCP Map

The Select MGCP Map dialog box lets you select or create a new MGCP map. An MGCP map lets you change the configuration values used for MGCP application inspection. The Select MGCP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default MGCP inspection map—Specifies to use the default MGCP map.

- Select an MGCP map for fine control over inspection—Lets you select a defined application inspection map or add a new one.

- Add—Opens the Add Policy Map dialog box for the inspection.

MGCP Inspect Map

Configuration > Global Objects > Inspect Maps > MGCP

The MGCP pane lets you view previously configured MGCP application inspection maps. An MGCP map lets you change the default configuration values used for MGCP application inspection. You can use an MGCP map to manage connections between VoIP devices and MGCP call agents.

Fields

- MGCP Inspect Maps—Table that lists the defined MGCP inspect maps.

- Add—Configures a new MGCP inspect map.

- Edit—Edits the selected MGCP entry in the MGCP Inspect Maps table.

- Delete—Deletes the inspect map selected in the MGCP Inspect Maps table.

Gateways and Call Agents

Configuration > Global Objects > Inspect Maps > MGCP > Gateways and Call Agents

The Gateways and Call Agents dialog box lets you configure groups of gateways and call agents for the map.

Fields

- Group ID—Identifies the ID of the call agent group. A call agent group associates one or more call agents with one or more MGCP media gateways. The gateway IP address can only be associated with one group ID. You cannot use the same gateway with different group IDs. The valid range is from 0 to 2147483647Criterion—Shows the criterion of the inspection.

- Gateways—Identifies the IP address of the media gateway that is controlled by the associated call agent. A media gateway is typically a network element that provides conversion between the audio signals carried on telephone circuits and data packets carried over the Internet or over other packet networks. Normally, a gateway sends commands to the default MGCP port for call agents, 2727.

- Call Agents—Identifies the IP address of a call agent that controls the MGCP media gateways in the call agent group. Normally, a call agent sends commands to the default MGCP port for gateways, 2427.

- Add—Displays the Add MGCP dialog box, which you can use to define a new application inspection map.

- Edit—Displays the Edit MGCP dialog box, which you can use to modify the application inspection map selected in the application inspection map table.

- Delete—Deletes the application inspection map selected in the application inspection map table.

Add/Edit MGCP Policy Map

Configuration > Global Objects > Inspect Maps > MGCP > MGCP Inspect Map > View

The Add/Edit MGCP Policy Map pane lets you configure the command queue, gateway, and call agent settings for MGCP application inspection maps.

Fields

- Name—When adding an MGCP map, enter the name of the MGCP map. When editing an MGCP map, the name of the previously configured MGCP map is shown.

- Description—Enter the description of the MGCP map, up to 200 characters in length.

- Command Queue—Tab that lets you specify the permitted queue size for MGCP commands.

  - Command Queue Size—Specifies the maximum number of commands to queue. The valid range is from 1 to 2147483647.

- Gateways and Call Agents—Tab that lets you configure groups of gateways and call agents for this map.

  - Group ID—Identifies the ID of the call agent group. A call agent group associates one or more call agents with one or more MGCP media gateways. The gateway IP address can only be associated with one group ID. You cannot use the same gateway with different group IDs. The valid range is from 0 to 2147483647Criterion—Shows the criterion of the inspection.

- Gateways—Identifies the IP address of the media gateway that is controlled by the associated call agent. A media gateway is typically a network element that provides conversion between the audio signals carried on telephone circuits and data packets carried over the Internet or over other packet networks. Normally, a gateway sends commands to the default MGCP port for call agents, 2727.

- Call Agents—Identifies the IP address of a call agent that controls the MGCP media gateways in the call agent group. Normally, a call agent sends commands to the default MGCP port for gateways, 2427.

- Add—Displays the Add MGCP Group dialog box, which you can use to define a new MGCP group of gateways and call agents.

- Edit—Displays the Edit MGCP dialog box, which you can use to modify the MGCP group selected in the Gateways and Call Agents table.

- Delete—Deletes the MGCP group selected in the Gateways and Call Agents table.

Add/Edit MGCP Group

Configuration > Global Objects > Inspect Maps > MGCP > Add/Edit MGCP Group

The Add/Edit MGCP Group dialog box lets you define the configuration of an MGCP group that will be used when MGCP application inspection is enabled.

Fields

- Group ID—Specifies the ID of the call agent group. A call agent group associates one or more call agents with one or more MGCP media gateways. The valid range is from 0 to 2147483647.

  - Gateway to Be Added—Specifies the IP address of the media gateway that is controlled by the associated call agent. A media gateway is typically a network element that provides conversion between the audio signals carried on telephone circuits and data packets carried over the Internet or over other packet networks. Normally, a gateway sends commands to the default MGCP port for call agents, 2727.

  - Add—Adds the specified IP address to the IP address table.

  - Delete—Deletes the selected IP address from the IP address table.

  - IP Address—Lists the IP addresses of the gateways in the call agent group.

- Call Agents

  - Call Agent to Be Added—Specifies the IP address of a call agent that controls the MGCP media gateways in the call agent group. Normally, a call agent sends commands to the default MGCP port for gateways, 2427.

  - Add—Adds the specified IP address to the IP address table.

  - Delete—Deletes the selected IP address from the IP address table.

  - IP Address—Lists the IP addresses of the call agents in the call agent group.

RTSP Inspection

This section describes RTSP application inspection. This section includes the following topics:

-

RTSP Inspection Overview

The RTSP inspection engine lets the ASA pass RTSP packets. RTSP is used by RealAudio, RealNetworks, Apple QuickTime 4, RealPlayer, and Cisco IP/TV connections.

Note For Cisco IP/TV, use RTSP TCP port 554 and TCP 8554.

RTSP applications use the well-known port 554 with TCP (rarely UDP) as a control channel. The ASA only supports TCP, in conformity with RFC 2326. This TCP control channel is used to negotiate the data channels that is used to transmit audio/video traffic, depending on the transport mode that is configured on the client.

The supported RDT transports are: rtp/avp, rtp/avp/udp, x-real-rdt, x-real-rdt/udp, and x-pn-tng/udp.

The ASA parses Setup response messages with a status code of 200. If the response message is travelling inbound, the server is outside relative to the ASA and dynamic channels need to be opened for connections coming inbound from the server. If the response message is outbound, then the ASA does not need to open dynamic channels.

Because RFC 2326 does not require that the client and server ports must be in the SETUP response message, the ASA keeps state and remembers the client ports in the SETUP message. QuickTime places the client ports in the SETUP message and then the server responds with only the server ports.

RTSP inspection does not support PAT or dual-NAT. Also, the ASA cannot recognize HTTP cloaking where RTSP messages are hidden in the HTTP messages.

Using RealPlayer

When using RealPlayer, it is important to properly configure transport mode. For the ASA, add an **access-list** command from the server to the client or vice versa. For RealPlayer, change transport mode by clicking **Options>Preferences>Transport>RTSP Settings**.

If using TCP mode on the RealPlayer, select the **Use TCP to Connect to Server** and **Attempt to use TCP for all content** check boxes. On the ASA, there is no need to configure the inspection engine.

If using UDP mode on the RealPlayer, select the **Use TCP to Connect to Server** and **Attempt to use UDP for static content** check boxes, and for live content not available via Multicast. On the ASA, add an **inspect rtsp** *port* command.

Restrictions and Limitations

The following restrictions apply to the RSTP inspection.

- The ASA does not support multicast RTSP or RTSP messages over UDP.

- The ASA does not have the ability to recognize HTTP cloaking where RTSP messages are hidden in the HTTP messages.

- The ASA cannot perform NAT on RTSP messages because the embedded IP addresses are contained in the SDP files as part of HTTP or RTSP messages. Packets could be fragmented and ASA cannot perform NAT on fragmented packets.

- With Cisco IP/TV, the number of translates the ASA performs on the SDP part of the message is proportional to the number of program listings in the Content Manager (each program listing can have at least six embedded IP addresses).

- You can configure NAT for Apple QuickTime 4 or RealPlayer. Cisco IP/TV only works with NAT if the Viewer and Content Manager are on the outside network and the server is on the inside network.

Select RTSP Map

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select NetBIOS Map

The Select RTSP Map dialog box lets you select or create a new RTSP map. An RTSP map lets you change the configuration values used for RTSP application inspection. The Select RTSP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default RTSP inspection map—Specifies to use the default RTSP inspection map.

- Select a RTSP inspect map for fine control over inspection—Lets you select a defined application inspection map or add a new one.

- Add—Opens the Add Policy Map dialog box for the inspection.

RTSP Inspect Map

Configuration > Global Objects > Inspect Maps > RADIUS

The RTSP pane lets you view previously configured RTSP application inspection maps. An RTSP map lets you change the default configuration values used for RTSP application inspection. You can use an RTSP map to protect RTSP traffic.

Fields

- RTSP Inspect Maps—Table that lists the defined RTSP inspect maps.

- Add—Configures a new RTSP inspect map.

- Edit—Edits the selected RTSP entry in the RTSP Inspect Maps table.

- Delete—Deletes the inspect map selected in the RTSP Inspect Maps table.

Add/Edit RTSP Policy Map

Configuration > Global Objects > Inspect Maps > MGCP > MGCP Inspect Map > View

The Add/Edit RTSP Policy Map pane lets you configure the parameters and inspections settings for RTSP application inspection maps.

Fields

- Name—When adding an RTSP map, enter the name of the RTSP map. When editing an RTSP map, the name of the previously configured RTSP map is shown.

- Description—Enter the description of the RTSP map, up to 200 characters in length.

- Parameters—Tab that lets you restrict usage on reserved ports during media port negotiation, and lets you set the URL length limit.

  - Enforce Reserve Port Protection—Lets you restrict the use of reserved ports during media port negotiation.

  - Maximum URL Length—Specifies the maximum length of the URL allowed in the message. Maximum value is 6000.

- Inspections—Tab that shows you the RTSP inspection configuration and lets you add or edit.

  - Match Type—Shows the match type, which can be a positive or negative match.

  - Criterion—Shows the criterion of the RTSP inspection.

  - Value—Shows the value to match in the RTSP inspection.

  - Action—Shows the action if the match condition is met.

  - Log—Shows the log state.

  - Add—Opens the Add RTSP Inspect dialog box to add a RTSP inspection.

  - Edit—Opens the Edit RTSP Inspect dialog box to edit a RTSP inspection.

  - Delete—Deletes a RTSP inspection.

  - Move Up—Moves an inspection up in the list.

  - Move Down—Moves an inspection down in the list.

RTSP Class Map

Configuration > Firewall > Objects > Class Maps > RTSP

The RTSP Class Map pane lets you configure RTSP class maps for RTSP inspection.

An inspection class map matches application traffic with criteria specific to the application. You then identify the class map in the inspect map and enable actions. The difference between creating a class map and defining the traffic match directly in the inspect map is that you can create more complex match criteria and you can reuse class maps. The applications that support inspection class maps are DNS, FTP, H.323, HTTP, IM, SIP, and RTSP.

Fields

- Name—Shows the RTSP class map name.

- Match Conditions—Shows the type, match criterion, and value in the class map.

  - Match Type—Shows the match type, which can be a positive or negative match.

- – Criterion—Shows the criterion of the RTSP class map.
- – Value—Shows the value to match in the RTSP class map.
- Description—Shows the description of the class map.
- Add—Adds a RTSP class map.
- Edit—Edits a RTSP class map.
- Delete—Deletes a RTSP class map.

Add/Edit RTSP Traffic Class Map

Configuration > Firewall > Objects > Class Maps > RTSP > Add/Edit RTSP Traffic Class Map

The Add/Edit RTSP Traffic Class Map dialog box lets you define the match criterion, values, and actions for the RTSP traffic class map.

Fields

- Match Type—Specifies whether traffic should match or not match the values.

  For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map.

- Criterion—Specifies which criterion of RTSP traffic to match.
  - – URL Filter—Match URL filtering.
  - – Request Method—Match an RTSP request method.

- URL Filter Criterion Values—Specifies to match URL filtering. Applies the regular expression match.
  - – Regular Expression—Lists the defined regular expressions to match.
  - – Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - – Regular Expression Class—Lists the defined regular expression classes to match.
  - – Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- URL Filter Actions—Primary action and log settings.
  - – Action—Drop connection or log.
  - – Log—Enable or disable.

- Request Method Criterion Values—Specifies to match an RTSP request method.
  - – Request Method—Specifies a request method: announce, describe, get_parameter, options, pause, play, record, redirect, setup, set_parameters, teardown.

- Request Method Actions—Primary action settings.
  - – Action—Limit rate (pps).

SIP Inspection

This section describes SIP application inspection. This section includes the following topics:

SIP Inspection Overview

SIP, as defined by the IETF, enables call handling sessions, particularly two-party audio conferences, or "calls." SIP works with SDP for call signalling. SDP specifies the ports for the media stream. Using SIP, the ASA can support any SIP VoIP gateways and VoIP proxy servers. SIP and SDP are defined in the following RFCs:

- SIP: Session Initiation Protocol, RFC 3261
- SDP: Session Description Protocol, RFC 2327

To support SIP calls through the ASA, signaling messages for the media connection addresses, media ports, and embryonic connections for the media must be inspected, because while the signaling is sent over a well-known destination port (UDP/TCP 5060), the media streams are dynamically allocated. Also, SIP embeds IP addresses in the user-data portion of the IP packet. SIP inspection applies NAT for these embedded IP addresses.

The following limitations and restrictions apply when using PAT with SIP:

- If a remote endpoint tries to register with a SIP proxy on a network protected by the ASA, the registration fails under very specific conditions, as follows:
  - PAT is configured for the remote endpoint.
  - The SIP registrar server is on the outside network.
  - The port is missing in the contact field in the REGISTER message sent by the endpoint to the proxy server.
  - Configuring static PAT is not supported with SIP inspection. If static PAT is configured for the Cisco Unified Communications Manager, SIP inspection cannot rewrite the SIP packet. Configure one-to-one static NAT for the Cisco Unified Communications Manager.

- If a SIP device transmits a packet in which the SDP portion has an IP address in the owner/creator field (o=) that is different than the IP address in the connection field (c=), the IP address in the o= field may not be properly translated. This is due to a limitation in the SIP protocol, which does not provide a port value in the o= field.

- When using PAT, any SIP header field which contains an internal IP address without a port might not be translated and hence the internal IP address will be leaked outside. If you want to avoid this leakage, configure NAT instead of PAT.

SIP Instant Messaging

Instant Messaging refers to the transfer of messages between users in near real-time. SIP supports the Chat feature on Windows XP using Windows Messenger RTC Client version 4.7.0105 only. The MESSAGE/INFO methods and 202 Accept response are used to support IM as defined in the following RFCs:

- Session Initiation Protocol (SIP)-Specific Event Notification, RFC 3265

- Session Initiation Protocol (SIP) Extension for Instant Messaging, RFC 3428

MESSAGE/INFO requests can come in at any time after registration/subscription. For example, two users can be online at any time, but not chat for hours. Therefore, the SIP inspection engine opens pinholes that time out according to the configured SIP timeout value. This value must be configured at least five minutes longer than the subscription duration. The subscription duration is defined in the Contact Expires value and is typically 30 minutes.

Because MESSAGE/INFO requests are typically sent using a dynamically allocated port other than port 5060, they are required to go through the SIP inspection engine.

Note   Only the Chat feature is currently supported. Whiteboard, File Transfer, and Application Sharing are not supported. RTC Client 5.0 is not supported.

SIP inspection translates the SIP text-based messages, recalculates the content length for the SDP portion of the message, and recalculates the packet length and checksum. It dynamically opens media connections for ports specified in the SDP portion of the SIP message as address/ports on which the endpoint should listen.

SIP inspection has a database with indices CALL_ID/FROM/TO from the SIP payload. These indices identify the call, the source, and the destination. This database contains the media addresses and media ports found in the SDP media information fields and the media type. There can be multiple media addresses and ports for a session. The ASA opens RTP/RTCP connections between the two endpoints using these media addresses/ports.

The well-known port 5060 must be used on the initial call setup (INVITE) message; however, subsequent messages may not have this port number. The SIP inspection engine opens signaling connection pinholes, and marks these connections as SIP connections. This is done for the messages to reach the SIP application and be translated.

As a call is set up, the SIP session is in the "transient" state until the media address and media port is received from the called endpoint in a Response message indicating the RTP port the called endpoint listens on. If there is a failure to receive the response messages within one minute, the signaling connection is torn down.

Once the final handshake is made, the call state is moved to active and the signaling connection remains until a BYE message is received.

If an inside endpoint initiates a call to an outside endpoint, a media hole is opened to the outside interface to allow RTP/RTCP UDP packets to flow to the inside endpoint media address and media port specified in the INVITE message from the inside endpoint. Unsolicited RTP/RTCP UDP packets to an inside interface does not traverse the ASA, unless the ASA configuration specifically allows it.

Select SIP Map

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select SIP Map

The Select SIP Map dialog box lets you select or create a new SIP map. A SIP map lets you change the configuration values used for SIP application inspection. The Select SIP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default SIP inspection map—Specifies to use the default SIP map.

- Select a SIP map for fine control over inspection—Lets you select a defined application inspection map or add a new one.

- Add—Opens the Add Policy Map dialog box for the inspection.

- Enable encrypted traffic inspection check box—Select to enable the radio buttons to select a proxy type.

- Proxy Type

  - TLS Proxy radio button—Use TLS Proxy to enable inspection of encrypted traffic.

  - Phone Proxy radio button—Specifies to associate the Phone Proxy with the TLS Proxy that you select from the TLS Proxy Name field.

    Configure button—Opens the Configure the Phone Proxy dialog box so that you can specify or edit Phone Proxy configuration settings.

  - UC-IME Proxy ratio button—Specifies to associate the UC-IME Proxy (Cisco Intercompany Media Engine proxy) with the TLS Proxy that you select from the TLS Proxy Name field.

    Configure button—Opens the Configure the UC-IME Proxy dialog box so that you can specify or edit UC-IME Proxy configuration settings.

- TLS Proxy Name:—Name of existing TLS Proxy.

- Manage—Opens the Add TLS Proxy dialog box to add a TLS Proxy.

Only one TLS proxy can be assigned to the Phone Proxy or UC-IME Proxy at a time. If you configure more than one service policy rule for Phone Proxy or UC-IME Proxy inspection and attempt to assign a different TLS proxy to them, ASDM displays a warning that all other service policy rules with Phone Proxy or UC-IME inspection will be changed to use the latest selected TLS proxy.

The UC-IME Proxy configuration requires two TLS proxies – one for outbound traffic and one for inbound. Rather than associating the TLS proxies directly with the UC-IME Proxy, as is the case with phone proxy, the TLS proxies are associated with it indirectly via SIP inspection rules.

You associate a TLS proxy with the Phone Proxy while defining a SIP inspection action . ASDM will convert the association to the existing phone proxy.

SIP Class Map

Configuration > Global Objects > Class Maps > SIP

The SIP Class Map pane lets you configure SIP class maps for SIP inspection.

An inspection class map matches application traffic with criteria specific to the application. You then identify the class map in the inspect map and enable actions. The difference between creating a class map and defining the traffic match directly in the inspect map is that you can create more complex match criteria and you can reuse class maps. The applications that support inspection class maps are DNS, FTP, H.323, HTTP, IM, and SIP.

Fields

- Name—Shows the SIP class map name.
- Match Conditions—Shows the type, match criterion, and value in the class map.
    - Match Type—Shows the match type, which can be a positive or negative match.
    - Criterion—Shows the criterion of the SIP class map.
    - Value—Shows the value to match in the SIP class map.
- Description—Shows the description of the class map.
- Add—Adds a SIP class map.
- Edit—Edits a SIP class map.
- Delete—Deletes a SIP class map.

Add/Edit SIP Traffic Class Map

Configuration > Global Objects > Class Maps > SIP > Add/Edit SIP Traffic Class Map

The Add/Edit SIP Traffic Class Map dialog box lets you define a SIP class map.

Fields

- Name—Enter the name of the SIP class map, up to 40 characters in length.
- Description—Enter the description of the SIP class map.
- Add—Adds a SIP class map.
- Edit—Edits a SIP class map.
- Delete—Deletes a SIP class map.

Add/Edit SIP Match Criterion

Configuration > Global Objects > Class Maps > SIP > Add/Edit SIP Traffic Class Map > Add/Edit SIP Match Criterion

The Add/Edit SIP Match Criterion dialog box lets you define the match criterion and value for the SIP class map.

Fields

- Match Type—Specifies whether the class map includes traffic that matches the criterion, or traffic that does not match the criterion.

    For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map.

- Criterion—Specifies which criterion of SIP traffic to match.
    - Called Party—Match the called party as specified in the To header.
    - Calling Party—Match the calling party as specified in the From header.
    - Content Length—Match the Content Length header, between 0 and 65536.
    - Content Type—Match the Content Type header.
    - IM Subscriber—Match the SIP IM subscriber.

- – Message Path—Match the SIP Via header.

- – Request Method—Match the SIP request method.

- – Third-Party Registration—Match the requester of a third-party registration.

- – URI Length—Match a URI in the SIP headers, between 0 and 65536.

- Called Party Criterion Values—Specifies to match the called party. Applies the regular expression match.

   - – Regular Expression—Lists the defined regular expressions to match.

   - – Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

   - – Regular Expression Class—Lists the defined regular expression classes to match.

   - – Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Calling Party Criterion Values—Specifies to match the calling party. Applies the regular expression match.

   - – Regular Expression—Lists the defined regular expressions to match.

   - – Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

   - – Regular Expression Class—Lists the defined regular expression classes to match.

   - – Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Content Length Criterion Values—Specifies to match a SIP content header of a length greater than specified.

   - – Greater Than Length—Enter a header length value in bytes.

- Content Type Criterion Values—Specifies to match a SIP content header type.

   - – SDP—Match an SDP SIP content header type.

   - – Regular Expression—Match a regular expression.

      Regular Expression—Lists the defined regular expressions to match.

      Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

      Regular Expression Class—Lists the defined regular expression classes to match.

      Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- IM Subscriber Criterion Values—Specifies to match the IM subscriber. Applies the regular expression match.

   - – Regular Expression—Lists the defined regular expressions to match.

   - – Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

   - – Regular Expression Class—Lists the defined regular expression classes to match.

   - – Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Message Path Criterion Values—Specifies to match a SIP Via header. Applies the regular expression match.

    – Regular Expression—Lists the defined regular expressions to match.

    – Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

    – Regular Expression Class—Lists the defined regular expression classes to match.

    – Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Request Method Criterion Values—Specifies to match a SIP request method.

    – Request Method—Specifies a request method: ack, bye, cancel, info, invite, message, notify, options, prack, refer, register, subscribe, unknown, update.

- Third-Party Registration Criterion Values—Specifies to match the requester of a third-party registration. Applies the regular expression match.

    – Regular Expression—Lists the defined regular expressions to match.

    – Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

    – Regular Expression Class—Lists the defined regular expression classes to match.

    – Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- URI Length Criterion Values—Specifies to match a URI of a selected type and greater than the specified length in the SIP headers.

    – URI type—Specifies to match either SIP URI or TEL URI.

    – Greater Than Length—Length in bytes.

SIP Inspect Map

Configuration > Global Objects > Inspect Maps > SIP

The SIP pane lets you view previously configured SIP application inspection maps. A SIP map lets you change the default configuration values used for SIP application inspection.

SIP is a widely used protocol for Internet conferencing, telephony, presence, events notification, and instant messaging. Partially because of its text-based nature and partially because of its flexibility, SIP networks are subject to a large number of security threats.

SIP application inspection provides address translation in message header and body, dynamic opening of ports and basic sanity checks. It also supports application security and protocol conformance, which enforce the sanity of the SIP messages, as well as detect SIP-based attacks.

Fields

- SIP Inspect Maps—Table that lists the defined SIP inspect maps.

- Add—Configures a new SIP inspect map. To edit a SIP inspect map, choose the SIP entry in the SIP Inspect Maps table and click **Customize**.

- Delete—Deletes the inspect map selected in the SIP Inspect Maps table.

- Security Level—Select the security level (high or low).

    – Low—Default.

SIP instant messaging (IM) extensions: Enabled.

Non-SIP traffic on SIP port: Permitted.

Hide server's and endpoint's IP addresses: Disabled.

Mask software version and non-SIP URIs: Disabled.

Ensure that the number of hops to destination is greater than 0: Enabled.

RTP conformance: Not enforced.

SIP conformance: Do not perform state checking and header validation.

– Medium

SIP instant messaging (IM) extensions: Enabled.

Non-SIP traffic on SIP port: Permitted.

Hide server's and endpoint's IP addresses: Disabled.

Mask software version and non-SIP URIs: Disabled.

Ensure that the number of hops to destination is greater than 0: Enabled.

RTP conformance: Enforced.

Limit payload to audio or video, based on the signaling exchange: No

SIP conformance: Drop packets that fail state checking.

– High

SIP instant messaging (IM) extensions: Enabled.

Non-SIP traffic on SIP port: Denied.

Hide server's and endpoint's IP addresses: Disabled.

Mask software version and non-SIP URIs: Enabled.

Ensure that the number of hops to destination is greater than 0: Enabled.

RTP conformance: Enforced.

Limit payload to audio or video, based on the signaling exchange: Yes

SIP conformance: Drop packets that fail state checking and packets that fail header validation.

– Customize—Opens the Add/Edit SIP Policy Map dialog box for additional settings.

– Default Level—Sets the security level back to the default level of Low.

Add/Edit SIP Policy Map (Security Level)

Configuration > Global Objects > Inspect Maps > SIP > SIP Inspect Map > Basic View

The Add/Edit SIP Policy Map pane lets you configure the security level and additional settings for SIP application inspection maps.

Fields

• Name—When adding a SIP, enter the name of the SIP map. When editing a SIP map, the name of the previously configured SIP map is shown.

• Description—Enter the description of the SIP map, up to 200 characters in length.

• Security Level—Select the security level (high or low).

    – Low—Default.

    SIP instant messaging (IM) extensions: Enabled.

    Non-SIP traffic on SIP port: Permitted.

    Hide server's and endpoint's IP addresses: Disabled.

    Mask software version and non-SIP URIs: Disabled.

    Ensure that the number of hops to destination is greater than 0: Enabled.

    RTP conformance: Not enforced.

    SIP conformance: Do not perform state checking and header validation.

    – Medium

    SIP instant messaging (IM) extensions: Enabled.

    Non-SIP traffic on SIP port: Permitted.

    Hide server's and endpoint's IP addresses: Disabled.

    Mask software version and non-SIP URIs: Disabled.

    Ensure that the number of hops to destination is greater than 0: Enabled.

    RTP conformance: Enforced.

    Limit payload to audio or video, based on the signaling exchange: No

    SIP conformance: Drop packets that fail state checking.

    – High

    SIP instant messaging (IM) extensions: Enabled.

    Non-SIP traffic on SIP port: Denied.

    Hide server's and endpoint's IP addresses: Disabled.

    Mask software version and non-SIP URIs: Enabled.

    Ensure that the number of hops to destination is greater than 0: Enabled.

    RTP conformance: Enforced.

    Limit payload to audio or video, based on the signaling exchange: Yes

    SIP conformance: Drop packets that fail state checking and packets that fail header validation.

    – Default Level—Sets the security level back to the default.

- Details—Shows additional filtering, IP address privacy, hop count, RTP conformance, SIP conformance, field masking, and inspections settings to configure.

Add/Edit SIP Policy Map (Details)

Configuration > Global Objects > Inspect Maps > SIP > SIP Inspect Map > Advanced View

The Add/Edit SIP Policy Map pane lets you configure the security level and additional settings for SIP application inspection maps.

Fields

- Name—When adding a SIP, enter the name of the SIP map. When editing a SIP map, the name of the previously configured SIP map is shown.

- Description—Enter the description of the SIP map, up to 200 characters in length.

- Security Level—Shows the security level settings to configure

- Filtering—Tab that lets you configure the filtering settings for SIP.

  - Enable SIP instant messaging (IM) extensions—Enables Instant Messaging extensions. Default is enabled.

  - Permit non-SIP traffic on SIP port—Permits non-SIP traffic on SIP port. Permitted by default.

- IP Address Privacy—Tab that lets you configure the IP address privacy settings for SIP.

  - Hide server's and endpoint's IP addresses—Enables IP address privacy. Disabled by default.

- Hop Count—Tab that lets you configure the hop count settings for SIP.

  - Ensure that number of hops to destination is greater than 0—Enables check for the value of Max-Forwards header is zero.

    Action—Drop packet, Drop Connection, Reset, Log.

    Log—Enable or Disable.

- RTP Conformance—Tab that lets you configure the RTP conformance settings for SIP.

  - Check RTP packets for protocol conformance—Checks RTP/RTCP packets flowing on the pinholes for protocol conformance.

    Limit payload to audio or video, based on the signaling exchange—Enforces payload type to be audio/video based on the signaling exchange.

- SIP Conformance—Tab that lets you configure the SIP conformance settings for SIP.

  - Enable state transition checking—Enables SIP state checking.

    Action—Drop packet, Drop Connection, Reset, Log.

    Log—Enable or Disable.

  - Enable strict validation of header fields—Enables validation of SIP header fields.

    Action—Drop packet, Drop Connection, Reset, Log.

    Log—Enable or Disable.

- Field Masking—Tab that lets you configure the field masking settings for SIP.

  - Inspect non-SIP URIs—Enables non-SIP URI inspection in Alert-Info and Call-Info headers.

    Action—Mask or Log.

    Log—Enable or Disable.

  - Inspect server's and endpoint's software version—Inspects SIP endpoint software version in User-Agent and Server headers.

    Action—Mask or Log.

    Log—Enable or Disable.

- Inspections—Tab that shows you the SIP inspection configuration and lets you add or edit.

  - Match Type—Shows the match type, which can be a positive or negative match.

  - Criterion—Shows the criterion of the SIP inspection.

  - Value—Shows the value to match in the SIP inspection.

  - Action—Shows the action if the match condition is met.

  - Log—Shows the log state.

- – Add—Opens the Add SIP Inspect dialog box to add a SIP inspection.
- – Edit—Opens the Edit SIP Inspect dialog box to edit a SIP inspection.
- – Delete—Deletes a SIP inspection.
- – Move Up—Moves an inspection up in the list.
- – Move Down—Moves an inspection down in the list.

Add/Edit SIP Inspect

Configuration > Global Objects > Inspect Maps > SIP > SIP Inspect Map > Advanced View > Add/Edit SIP Inspect

The Add/Edit SIP Inspect dialog box lets you define the match criterion and value for the SIP inspect map.

Fields

- Single Match—Specifies that the SIP inspect has only one match statement.
- Match Type—Specifies whether traffic should match or not match the values.

  For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map.

- Criterion—Specifies which criterion of SIP traffic to match.
    - – Called Party—Match a called party as specified in the To header.
    - – Calling Party—Match a calling party as specified in the From header.
    - – Content Length—Match a content length header.
    - – Content Type—Match a content type header.
    - – IM Subscriber—Match a SIP IM subscriber.
    - – Message Path—Match a SIP Via header.
    - – Request Method—Match a SIP request method.
    - – Third-Party Registration—Match the requester of a third-party registration.
    - – URI Length—Match a URI in the SIP headers.
- Called Party Criterion Values—Specifies to match the called party. Applies the regular expression match.
    - – Regular Expression—Lists the defined regular expressions to match.
    - – Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
    - – Regular Expression Class—Lists the defined regular expression classes to match.
    - – Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Calling Party Criterion Values—Specifies to match the calling party. Applies the regular expression match.
    - – Regular Expression—Lists the defined regular expressions to match.
    - – Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

- – Regular Expression Class—Lists the defined regular expression classes to match.

- – Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Content Length Criterion Values—Specifies to match a SIP content header of a length greater than specified.

  - – Greater Than Length—Enter a header length value in bytes.

- Content Type Criterion Values—Specifies to match a SIP content header type.

  - – SDP—Match an SDP SIP content header type.

  - – Regular Expression—Match a regular expression.

    Regular Expression—Lists the defined regular expressions to match.

    Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

    Regular Expression Class—Lists the defined regular expression classes to match.

    Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- IM Subscriber Criterion Values—Specifies to match the IM subscriber. Applies the regular expression match.

  - – Regular Expression—Lists the defined regular expressions to match.

  - – Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  - – Regular Expression Class—Lists the defined regular expression classes to match.

  - – Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Message Path Criterion Values—Specifies to match a SIP Via header. Applies the regular expression match.

  - – Regular Expression—Lists the defined regular expressions to match.

  - – Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  - – Regular Expression Class—Lists the defined regular expression classes to match.

  - – Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- Request Method Criterion Values—Specifies to match a SIP request method.

  - – Request Method—Specifies a request method: ack, bye, cancel, info, invite, message, notify, options, prack, refer, register, subscribe, unknown, update.

- Third-Party Registration Criterion Values—Specifies to match the requester of a third-party registration. Applies the regular expression match.

  - – Regular Expression—Lists the defined regular expressions to match.

  - – Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.

  - – Regular Expression Class—Lists the defined regular expression classes to match.

  - – Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.

- URI Length Criterion Values—Specifies to match a URI in the SIP headers greater than specified length.

    - URI type—Specifies to match either SIP URI or TEL URI.

    - Greater Than Length—Length in bytes.

- Multiple Matches—Specifies multiple matches for the SIP inspection.

    - SIP Traffic Class—Specifies the SIP traffic class match.

    - Manage—Opens the Manage SIP Class Maps dialog box to add, edit, or delete SIP Class Maps.

- Actions—Primary action and log settings.

    - Action—Drop packet, drop connection, reset, log. Note: Limit rate (pps) action is available for request methods invite and register.

    - Log—Enable or disable.

Skinny (SCCP) Inspection

This section describes SCCP application inspection. This section includes the following topics:

SCCP Inspection Overview

Note   For specific information about setting up the Phone Proxy on the ASA, which is part of the Cisco Unified Communications architecture and supports IP phone deployment, see Chapter 15, "Cisco Phone Proxy.".

Skinny (SCCP) is a simplified protocol used in VoIP networks. Cisco IP Phones using SCCP can coexist in an H.323 environment. When used with Cisco CallManager, the SCCP client can interoperate with H.323 compliant terminals.

The ASA supports PAT and NAT for SCCP. PAT is necessary if you have more IP phones than global IP addresses for the IP phones to use. By supporting NAT and PAT of SCCP Signaling packets, Skinny application inspection ensures that all SCCP signalling and media packets can traverse the ASA.

Normal traffic between Cisco CallManager and Cisco IP Phones uses SCCP and is handled by SCCP inspection without any special configuration. The ASA also supports DHCP options 150 and 66, which it accomplishes by sending the location of a TFTP server to Cisco IP Phones and other DHCP clients. Cisco IP Phones might also include DHCP option 3 in their requests, which sets the default route.

> **Note** The ASA supports inspection of traffic from Cisco IP Phones running SCCP protocol version 19 and earlier.

Supporting Cisco IP Phones

> **Note** For specific information about setting up the Phone Proxy on the ASA, which is part of the Cisco Unified Communications architecture and supports IP phone deployment, see Chapter 15, "Cisco Phone Proxy."

In topologies where Cisco CallManager is located on the higher security interface with respect to the Cisco IP Phones, if NAT is required for the Cisco CallManager IP address, the mapping must be **static** as a Cisco IP Phone requires the Cisco CallManager IP address to be specified explicitly in its configuration. An static identity entry allows the Cisco CallManager on the higher security interface to accept registrations from the Cisco IP Phones.

Cisco IP Phones require access to a TFTP server to download the configuration information they need to connect to the Cisco CallManager server.

When the Cisco IP Phones are on a lower security interface compared to the TFTP server, you must use an ACL to connect to the protected TFTP server on UDP port 69. While you do need a static entry for the TFTP server, this does not have to be an identity static entry. When using NAT, an identity static entry maps to the same IP address. When using PAT, it maps to the same IP address and port.

When the Cisco IP Phones are on a *higher* security interface compared to the TFTP server and Cisco CallManager, no ACL or static entry is required to allow the Cisco IP Phones to initiate the connection.

Restrictions and Limitations

The following are limitations that apply to the current version of PAT and NAT support for SCCP:

- PAT does not work with configurations containing the **alias** command.
- Outside NAT or PAT is *not* supported.

If the address of an internal Cisco CallManager is configured for NAT or PAT to a different IP address or port, registrations for external Cisco IP Phones fail because the ASA currently does not support NAT or PAT for the file content transferred over TFTP. Although the ASA supports NAT of TFTP messages and opens a pinhole for the TFTP file, the ASA cannot translate the Cisco CallManager IP address and port embedded in the Cisco IP Phone configuration files that are transferred by TFTP during phone registration.

> **Note** The ASA supports stateful failover of SCCP calls except for calls that are in the middle of call setup.

Select SCCP (Skinny) Map

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab >
Select SCCP Map

The Select SCCP (Skinny) Map dialog box lets you select or create a new SCCP (Skinny) map. An SCCP (Skinny) map lets you change the configuration values used for SCCP (Skinny) application inspection. The Select SCCP (Skinny) Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default SCCP (Skinny) inspection map—Specifies to use the default SCCP (Skinny) map.

- Select an SCCP (Skinny) map for fine control over inspection—Lets you select a defined application inspection map or add a new one.

- Add—Opens the Add Policy Map dialog box for the inspection.

- Encrypted Traffic Inspection—Lets you specify TLS proxy settings for the inspect map.

  – Do not inspect Encrypted Traffic—Disables the inspection of Skinny application inspection.

  – Use Phone Proxy to enable inspection of encrypted traffic—Uses the Phone Proxy configured on the ASA to inspect Skinny application traffic. See Chapter 15, "Cisco Phone Proxy."

  – Use TLS Proxy to enable inspection of encrypted traffic—Specifies to use Transaction Layer Security Proxy to enable inspection of encryped traffic.

    TLS Proxy Name:—Name of existing TLS Proxy.

    New—Opens the Add TLS Proxy dialog box to add a TLS Proxy.

SCCP (Skinny) Inspect Map

Configuration > Global Objects > Inspect Maps > SCCP (Skinny)

The SCCP (Skinny) pane lets you view previously configured SCCP (Skinny) application inspection maps. An SCCP (Skinny) map lets you change the default configuration values used for SCCP (Skinny) application inspection.

Skinny application inspection performs translation of embedded IP address and port numbers within the packet data, and dynamic opening of pinholes. It also performs additional protocol conformance checks and basic state tracking.

Fields

- SCCP (Skinny) Inspect Maps—Table that lists the defined SCCP (Skinny) inspect maps.

- Add—Configures a new SCCP (Skinny) inspect map. To edit an SCCP (Skinny) inspect map, choose the SCCP (Skinny) entry in the SCCP (Skinny) Inspect Maps table and click **Customize**.

- Delete—Deletes the inspect map selected in the SCCP (Skinny) Inspect Maps table.

- Security Level—Select the security level (high or low).

  – Low—Default.

    Registration: Not enforced.

    Maximum message ID: 0x181.

    Minimum prefix length: 4

    Media timeout: 00:05:00

    Signaling timeout: 01:00:00.

    RTP conformance: Not enforced.

– Medium

Registration: Not enforced.

Maximum message ID: 0x141.

Minimum prefix length: 4.

Media timeout: 00:01:00.

Signaling timeout: 00:05:00.

RTP conformance: Enforced.

Limit payload to audio or video, based on the signaling exchange: No.

– High

Registration: Enforced.

Maximum message ID: 0x141.

Minimum prefix length: 4.

Maximum prefix length: 65536.

Media timeout: 00:01:00.

Signaling timeout: 00:05:00.

RTP conformance: Enforced.

Limit payload to audio or video, based on the signaling exchange: Yes.

– Message ID Filtering—Opens the Messaging ID Filtering dialog box for configuring message ID filters.

– Customize—Opens the Add/Edit SCCP (Skinny) Policy Map dialog box for additional settings.

– Default Level—Sets the security level back to the default level of Low.

Message ID Filtering

Configuration > Global Objects > Inspect Maps > SCCP (Skinny) > Message ID Filtering

The Message ID Filtering dialog box lets you configure the settings for a message ID filter.

Fields

- Match Type—Shows the match type, which can be a positive or negative match.
- Criterion—Shows the criterion of the inspection.
- Value—Shows the value to match in the inspection.
- Action—Shows the action if the match condition is met.
- Log—Shows the log state.
- Add—Opens the Add Message ID Filtering dialog box to add a message ID filter.
- Edit—Opens the Edit Message ID Filtering dialog box to edit a message ID filter.
- Delete—Deletes a message ID filter.
- Move Up—Moves an entry up in the list.
- Move Down—Moves an entry down in the list.

Add/Edit SCCP (Skinny) Policy Map (Security Level)

Configuration > Global Objects > Inspect Maps > SCCP (Skinny) > SCCP (Skinny) Inspect Map > Basic View

The Add/Edit SCCP (Skinny) Policy Map pane lets you configure the security level and additional settings for SCCP (Skinny) application inspection maps.

Fields

- Name—When adding an SCCP (Skinny) map, enter the name of the SCCP (Skinny) map. When editing an SCCP (Skinny) map, the name of the previously configured SCCP (Skinny) map is shown.

- Description—Enter the description of the SCCP (Skinny) map, up to 200 characters in length.

- Security Level—Select the security level (high or low).

  - Low—Default.

    Registration: Not enforced.

    Maximum message ID: 0x181.

    Minimum prefix length: 4

    Media timeout: 00:05:00

    Signaling timeout: 01:00:00.

    RTP conformance: Not enforced.

  - Medium

    Registration: Not enforced.

    Maximum message ID: 0x141.

    Minimum prefix length: 4.

    Media timeout: 00:01:00.

    Signaling timeout: 00:05:00.

    RTP conformance: Enforced.

    Limit payload to audio or video, based on the signaling exchange: No.

  - High

    Registration: Enforced.

    Maximum message ID: 0x141.

    Minimum prefix length: 4.

    Maximum prefix length: 65536.

    Media timeout: 00:01:00.

    Signaling timeout: 00:05:00.

    RTP conformance: Enforced.

    Limit payload to audio or video, based on the signaling exchange: Yes.

  - Message ID Filtering—Opens the Messaging ID Filtering dialog box for configuring message ID filters.

  - Default Level—Sets the security level back to the default.

- Details—Shows additional parameter, RTP conformance, and message ID filtering settings to configure.

Add/Edit SCCP (Skinny) Policy Map (Details)

Configuration > Global Objects > Inspect Maps > SCCP (Skinny) > SCCP (Skinny) Inspect Map > Advanced View

The Add/Edit SCCP (Skinny) Policy Map pane lets you configure the security level and additional settings for SCCP (Skinny) application inspection maps.

Fields

- Name—When adding an SCCP (Skinny) map, enter the name of the SCCP (Skinny) map. When editing an SCCP (Skinny) map, the name of the previously configured SCCP (Skinny) map is shown.

- Description—Enter the description of the DNS map, up to 200 characters in length.

- Security Level—Shows the security level and message ID filtering settings to configure.

- Parameters—Tab that lets you configure the parameter settings for SCCP (Skinny).

    - Enforce endpoint registration—Enforce that Skinny endpoints are registered before placing or receiving calls.

        Maximum Message ID—Specify value of maximum SCCP message ID allowed.

    - SCCP Prefix Length—Specifies prefix length value in Skinny messages.

        Minimum Prefix Length—Specify minimum value of SCCP prefix length allowed.

        Maximum Prefix Length—Specify maximum value of SCCP prefix length allowed.

    - Media Timeout—Specify timeout value for media connections.

    - Signaling Timeout—Specify timeout value for signaling connections.

- RTP Conformance—Tab that lets you configure the RTP conformance settings for SCCP (Skinny).

    - Check RTP packets for protocol conformance—Checks RTP/RTCP packets flowing on the pinholes for protocol conformance.

        Limit payload to audio or video, based on the signaling exchange—Enforces the payload type to be audio/video based on the signaling exchange.

- Message ID Filtering—Tab that lets you configure the message ID filtering settings for SCCP (Skinny).

    - Match Type—Shows the match type, which can be a positive or negative match.

    - Criterion—Shows the criterion of the inspection.

    - Value—Shows the value to match in the inspection.

    - Action—Shows the action if the match condition is met.

    - Log—Shows the log state.

    - Add—Opens the Add Message ID Filtering dialog box to add a message ID filter.

    - Edit—Opens the Edit Message ID Filtering dialog box to edit a message ID filter.

    - Delete—Deletes a message ID filter.

    - Move Up—Moves an entry up in the list.

    - Move Down—Moves an entry down in the list.

Add/Edit Message ID Filter

Configuration > Global Objects > Inspect Maps > SCCP (Skinny) > SCCP (Skinny) Inspect Map > Advanced View > Add/Edit Message ID Filter

The Add Message ID Filter dialog box lets you configure message ID filters.

Fields

- Match Type—Specifies whether traffic should match or not match the values.

    For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map.

- Criterion—Specifies which criterion of SCCP (Skinny) traffic to match.

    – Message ID—Match specified message ID.

    Message ID—Specify value of maximum SCCP message ID allowed.

    – Message ID Range—Match specified message ID range.

    Lower Message ID—Specify lower value of SCCP message ID allowed.

    Upper Message ID—Specify upper value of SCCP message ID allowed.

- Action—Drop packet.

- Log—Enable or disable.

Inspection of Database and Directory Protocols

This chapter describes how to configure application layer protocol inspection. Inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the ASA to do a deep packet inspection instead of passing the packet through the fast path. As a result, inspection engines can affect overall throughput.

Several common inspection engines are enabled on the ASA by default, but you might need to enable others depending on your network.

This chapter includes the following sections:

- ILS Inspection, page 11-1
- SQL*Net Inspection, page 11-2
- Sun RPC Inspection, page 11-3

ILS Inspection

The ILS inspection engine provides NAT support for Microsoft NetMeeting, SiteServer, and Active Directory products that use LDAP to exchange directory information with an ILS server.

The ASA supports NAT for ILS, which is used to register and locate endpoints in the ILS or SiteServer Directory. PAT cannot be supported because only IP addresses are stored by an LDAP database.

For search responses, when the LDAP server is located outside, NAT should be considered to allow internal peers to communicate locally while registered to external LDAP servers. For such search responses, xlates are searched first, and then DNAT entries to obtain the correct address. If both of these searches fail, then the address is not changed. For sites using NAT 0 (no NAT) and not expecting DNAT interaction, we recommend that the inspection engine be turned off to provide better performance.

Additional configuration may be necessary when the ILS server is located inside the ASA border. This would require a hole for outside clients to access the LDAP server on the specified port, typically TCP 389.

Because ILS traffic only occurs on the secondary UDP channel, the TCP connection is disconnected after the TCP inactivity interval. By default, this interval is 60 minutes and can be adjusted using the **timeout** command.

ILS/LDAP follows a client/server model with sessions handled over a single TCP connection. Depending on the client's actions, several of these sessions may be created.

During connection negotiation time, a BIND PDU is sent from the client to the server. Once a successful BIND RESPONSE from the server is received, other operational messages may be exchanged (such as ADD, DEL, SEARCH, or MODIFY) to perform operations on the ILS Directory. The ADD REQUEST and SEARCH RESPONSE PDUs may contain IP addresses of NetMeeting peers, used by H.323 (SETUP and CONNECT messages) to establish the NetMeeting sessions. Microsoft NetMeeting v2.X and v3.X provides ILS support.

The ILS inspection performs the following operations:

- Decodes the LDAP REQUEST/RESPONSE PDUs using the BER decode functions
- Parses the LDAP packet
- Extracts IP addresses
- Translates IP addresses as necessary
- Encodes the PDU with translated addresses using BER encode functions
- Copies the newly encoded PDU back to the TCP packet
- Performs incremental TCP checksum and sequence number adjustment

ILS inspection has the following limitations:

- Referral requests and responses are not supported
- Users in multiple directories are not unified
- Single users having multiple identities in multiple directories cannot be recognized by NAT

Note   Because H.225 call signalling traffic only occurs on the secondary UDP channel, the TCP connection is disconnected after the interval specified by the TCP option in the Configuration > Firewall > Advanced > Global Timeouts pane. By default, this interval is set at 60 minutes.

SQL*Net Inspection

SQL*Net inspection is enabled by default.

The SQL*Net protocol consists of different packet types that the ASA handles to make the data stream appear consistent to the Oracle applications on either side of the ASA.

The default port assignment for SQL*Net is 1521. This is the value used by Oracle for SQL*Net, but this value does not agree with IANA port assignments for Structured Query Language (SQL).

Note   Disable SQL*Net inspection when SQL data transfer occurs on the same port as the SQL control TCP port 1521. The security appliance acts as a proxy when SQL*Net inspection is enabled and reduces the client window size from 65000 to about 16000 causing data transfer issues.

The ASA translates all addresses and looks in the packets for all embedded ports to open for SQL*Net Version 1.

For SQL*Net Version 2, all DATA or REDIRECT packets that immediately follow REDIRECT packets with a zero data length will be fixed up.

The packets that need fix-up contain embedded host/port addresses in the following format:

`(ADDRESS=(PROTOCOL=tcp)(DEV=6)(HOST=a.b.c.d)(PORT=a))`

SQL*Net Version 2 TNSFrame types (Connect, Accept, Refuse, Resend, and Marker) will not be scanned for addresses to NAT nor will inspection open dynamic connections for any embedded ports in the packet.

SQL*Net Version 2 TNSFrames, Redirect, and Data packets will be scanned for ports to open and addresses to NAT, if preceded by a REDIRECT TNSFrame type with a zero data length for the payload. When the Redirect message with data length zero passes through the ASA, a flag will be set in the connection data structure to expect the Data or Redirect message that follows to be translated and ports to be dynamically opened. If one of the TNS frames in the preceding paragraph arrive after the Redirect message, the flag will be reset.

The SQL*Net inspection engine will recalculate the checksum, change IP, TCP lengths, and readjust Sequence Numbers and Acknowledgment Numbers using the delta of the length of the new and old message.

SQL*Net Version 1 is assumed for all other cases. TNSFrame types (Connect, Accept, Refuse, Resend, Marker, Redirect, and Data) and all packets will be scanned for ports and addresses. Addresses will be translated and port connections will be opened.

Sun RPC Inspection

This section describes Sun RPC application inspection. This section includes the following topics:

Sun RPC Inspection Overview

The Sun RPC inspection engine enables or disables application inspection for the Sun RPC protocol. Sun RPC is used by NFS and NIS. Sun RPC services can run on any port. When a client attempts to access an Sun RPC service on a server, it must learn the port that service is running on. It does this by querying the port mapper process, usually rpcbind, on the well-known port of 111.

The client sends the Sun RPC program number of the service and the port mapper process responds with the port number of the service. The client sends its Sun RPC queries to the server, specifying the port identified by the port mapper process. When the server replies, the ASA intercepts this packet and opens both embryonic TCP and UDP connections on that port.

The following limitations apply to Sun RPC inspection:

- NAT or PAT of Sun RPC payload information is not supported.
- Sun RPC inspection supports inbound ACLs only. Sun RPC inspection does not support outbound ACLs because the inspection engine uses dynamic ACLs instead of secondary connections. Dynamic ACLs are always added on the ingress direction and not on egress; therefore, this inspection engine does not support outbound ACLs. To view the dynamic ACLs configured for the ASA, use the **show asp table classify domain permit** command. For information about the **show asp table classify domain permit** command, see the CLI configuration guide.

SUNRPC Server

Configuration > Properties > SUNRPC Server

The Configuration > Firewall > Advanced > SUNRPC Server pane shows which SunRPC services can traverse the ASA and their specific timeout, on a per server basis.

Fields

- Interface—Displays the interface on which the SunRPC server resides.

- IP address—Displays the IP address of the SunRPC server.

- Mask—Displays the subnet mask of the IP Address of the SunRPC server.

- Service ID—Displays the SunRPC program number, or service ID, allowed to traverse the ASA.

- Protocol—Displays the SunRPC transport protocol (TCP or UDP).

- Port—Displays the SunRPC protocol port range.

- Timeout—Displays the idle time after which the access for the SunRPC service traffic is closed.

Add/Edit SUNRPC Service

Configuration > Properties > SUNRPC Server > Add/Edit SUNRPC Service

The Configuration > Firewall > Advanced > SUNRPC Server > Add/Edit SUNRPC Service dialog box lets you specify what SunRPC services are allowed to traverse the ASA and their specific timeout, on a per-server basis.

Fields

- Interface Name—Specifies the interface on which the SunRPC server resides.

- Protocol—Specifies the SunRPC transport protocol (TCP or UDP).

- IP address—Specifies the IP address of the SunRPC server.

- Port—Specifies the SunRPC protocol port range.

- Mask—Specifies the subnet mask of the IP Address of the SunRPC server.

- Timeout—Specifies the idle time after which the access for the SunRPC service traffic is closed. Format is HH:MM:SS.

- Service ID—Specifies the SunRPC program number, or service ID, allowed to traverse the ASA.

CHAPTER **12**

Inspection for Management Application Protocols

This chapter describes how to configure application layer protocol inspection. Inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the ASA to do a deep packet inspection instead of passing the packet through the fast path. As a result, inspection engines can affect overall throughput.

Several common inspection engines are enabled on the ASA by default, but you might need to enable others depending on your network.

This chapter includes the following sections:

- DCERPC Inspection, page 12-1
- GTP Inspection, page 12-4
- RADIUS Accounting Inspection, page 12-10
- RSH Inspection, page 12-13
- SNMP Inspection, page 12-13
- XDMCP Inspection, page 12-15

DCERPC Inspection

This section describes the DCERPC inspection engine. This section includes the following topics:

- DCERPC Overview, page 12-1
- Select DCERPC Map, page 12-2
- DCERPC Inspect Map, page 12-2
- Add/Edit DCERPC Policy Map, page 12-3

DCERPC Overview

DCERPC is a protocol widely used by Microsoft distributed client and server applications that allows software clients to execute programs on a server remotely.

This typically involves a client querying a server called the Endpoint Mapper listening on a well known port number for the dynamically allocated network information of a required service. The client then sets up a secondary connection to the server instance providing the service. The security appliance allows the appropriate port number and network address and also applies NAT, if needed, for the secondary connection.

DCERPC inspect maps inspect for native TCP communication between the EPM and client on well known TCP port 135. Map and lookup operations of the EPM are supported for clients. Client and server can be located in any security zone. The embedded server IP address and Port number are received from the applicable EPM response messages. Since a client may attempt multiple connections to the server port returned by EPM, multiple use of pinholes are allowed, which have user configurable timeouts.

Note   DCERPC inspection only supports communication between the EPM and clients to open pinholes through theASA. Clients using RPC communication that does not use the EPM is not supported with DCERPC inspection.

Select DCERPC Map

**Add/Edit Service Policy Rule Wizard > Rule Actions >
Protocol Inspection Tab >Select DCERPC Map**

The Select DCERPC Map dialog box lets you select or create a new DCERPC map. A DCERPC map lets you change the configuration values used for DCERPC application inspection. The Select DCERPC Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default DCERPC inspection map—Specifies to use the default DCERPC map.

- Select a DCERPC map for fine control over inspection—Lets you select a defined application inspection map or add a new one.

- Add—Opens the Add Policy Map dialog box for the inspection.

DCERPC Inspect Map

Configuration > Global Objects > Inspect Maps > DCERPC

The DCERPC pane lets you view previously configured DCERPC application inspection maps. A DCERPC map lets you change the default configuration values used for DCERPC application inspection.

DCERPC is a protocol widely used by Microsoft distributed client and server applications that allows software clients to execute programs on a server remotely.

This typically involves a client querying a server called the Endpoint Mapper (EPM) listening on a well known port number for the dynamically allocated network information of a required service. The client then sets up a secondary connection to the server instance providing the service. The security appliance allows the appropriate port number and network address and also applies NAT, if needed, for the secondary connection.

DCERPC inspect maps inspect for native TCP communication between the EPM and client on well known TCP port 135. Map and lookup operations of the EPM are supported for clients. Client and server can be located in any security zone. The embedded server IP address and Port number are received from the applicable EPM response messages. Because a client may attempt multiple connections to the server port returned by EPM, multiple use of pinholes are allowed, which have user configurable timeouts.

Fields

- DCERPC Inspect Maps—Table that lists the defined DCERPC inspect maps.
- Add—Configures a new DCERPC inspect map. To edit a DCERPC inspect map, choose the DCERPC entry in the DCERPC Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the DCERPC Inspect Maps table.
- Security Level—Select the security level (high, medium, or low).
  - Low

    Pinhole timeout: 00:02:00

    Endpoint mapper service: not enforced

    Endpoint mapper service lookup: enabled

    Endpoint mapper service lookup timeout: 00:05:00
  - Medium—Default.

    Pinhole timeout: 00:01:00

    Endpoint mapper service: not enforced

    Endpoint mapper service lookup: disabled.
  - High

    Pinhole timeout: 00:01:00

    Endpoint mapper service: enforced

    Endpoint mapper service lookup: disabled
  - Customize—Opens the Add/Edit DCERPC Policy Map dialog box for additional settings.
  - Default Level—Sets the security level back to the default level of Medium.

Add/Edit DCERPC Policy Map

Configuration > Global Objects > Inspect Maps > DCERPC > DCERPC Inspect Map > Basic/Advanced View

The Add/Edit DCERPC Policy Map pane lets you configure the security level and parameters for DCERPC application inspection maps.

Fields

- Name—When adding a DCERPC map, enter the name of the DCERPC map. When editing a DCERPC map, the name of the previously configured DCERPC map is shown.
- Description—Enter the description of the DCERPC map, up to 200 characters in length.
- Security Level—Select the security level (high, medium, or low).
  - Low

    Pinhole timeout: 00:02:00

Endpoint mapper service: not enforced

Endpoint mapper service lookup: enabled

Endpoint mapper service lookup timeout: 00:05:00

– Medium—Default.

Pinhole timeout: 00:01:00

Endpoint mapper service: not enforced

Endpoint mapper service lookup: disabled.

– High

Pinhole timeout: 00:01:00

Endpoint mapper service: enforced

Endpoint mapper service lookup: disabled

– Default Level—Sets the security level back to the default level of Medium.

- Details—Shows the Parameters to configure additional settings.

– Pinhole Timeout—Sets the pinhole timeout. Because a client may use the server information returned by the endpoint mapper for multiple connections, the timeout value is configurable based on the client application environment. Range is from 0:0:1 to 1193:0:0. Default is 2 minutes.

– Enforce endpoint-mapper service—Enforces endpoint mapper service during binding.

– Enable endpoint-mapper service lookup—Enables the lookup operation of the endpoint mapper service. If disabled, the pinhole timeout is used.

Enforce Service Lookup Timeout—Enforces the service lookup timeout specified.

Service Lookup Timeout—Sets the timeout for pinholes from lookup operation.

GTP Inspection

This section describes the GTP inspection engine. This section includes the following topics:

Note GTP inspection requires a special license.

GTP Inspection Overview

GPRS provides uninterrupted connectivity for mobile subscribers between GSM networks and corporate networks or the Internet. The GGSN is the interface between the GPRS wireless data network and other networks. The SGSN performs mobility, data session management, and data compression (See Figure 12-1).

Figure 12-1   GPRS Tunneling Protocol

The UMTS is the commercial convergence of fixed-line telephony, mobile, Internet and computer technology. UTRAN is the networking protocol used for implementing wireless networks in this system. GTP allows multi-protocol packets to be tunneled through a UMTS/GPRS backbone between a GGSN, an SGSN and the UTRAN.

GTP does not include any inherent security or encryption of user data, but using GTP with the ASA helps protect your network against these risks.

The SGSN is logically connected to a GGSN using GTP. GTP allows multiprotocol packets to be tunneled through the GPRS backbone between GSNs. GTP provides a tunnel control and management protocol that allows the SGSN to provide GPRS network access for a mobile station by creating, modifying, and deleting tunnels. GTP uses a tunneling mechanism to provide a service for carrying user data packets.

Note   When using GTP with failover, if a GTP connection is established and the active unit fails before data is transmitted over the tunnel, the GTP data connection (with a "j" flag set) is not replicated to the standby unit. This occurs because the active unit does not replicate embryonic connections to the standby unit.

Select GTP Map

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab >Select GTP Map

The Select GTP Map dialog box lets you select or create a new GTP map. A GTP map lets you change the configuration values used for GTP application inspection. The Select GTP Map table provides a list of previously configured maps that you can select for application inspection.

Note   GTP inspection requires a special license. If you try to enable GTP application inspection on a ASA without the required license, the ASA displays an error message.

Fields

- Use the default GTP inspection map—Specifies to use the default GTP map.

- Select an GTP map for fine control over inspection—Lets you select a defined application inspection map or add a new one.

- Add—Opens the Add Policy Map dialog box for the inspection.

GTP Inspect Map

Configuration > Global Objects > Inspect Maps > GTP

The GTP pane lets you view previously configured GTP application inspection maps. A GTP map lets you change the default configuration values used for GTP application inspection.

GTP is a relatively new protocol designed to provide security for wireless connections to TCP/IP networks, such as the Internet. You can use a GTP map to control timeout values, message sizes, tunnel counts, and GTP versions traversing the security appliance.

Note   GTP inspection is not available without a special license.

Fields

- GTP Inspect Maps—Table that lists the defined GTP inspect maps.

- Add—Configures a new GTP inspect map. To edit a GTP inspect map, choose the GTP entry in the GTP Inspect Maps table and click **Customize**.

- Delete—Deletes the inspect map selected in the GTP Inspect Maps table.

- Security Level—Security level low only.

    - Do not Permit Errors

    - Maximum Number of Tunnels: 500

    - GSN timeout: 00:30:00

    - Pdp-Context timeout: 00:30:00

    - Request timeout: 00:01:00

    - Signaling timeout: 00:30:00.

    - Tunnel timeout: 01:00:00.

    - T3-response timeout: 00:00:20.

    - Drop and log unknown message IDs.

- IMSI Prefix Filtering—Opens the IMSI Prefix Filtering dialog box to configure IMSI prefix filters.

- Customize—Opens the Add/Edit GTP Policy Map dialog box for additional settings.

- Default Level—Sets the security level back to the default.

IMSI Prefix Filtering

Configuration > Global Objects > Inspect Maps > GTP > IMSI Prefix Filtering

The IMSI Prefix tab lets you define the IMSI prefix to allow within GTP requests.

Fields

- Mobile Country Code—Defines the non-zero, three-digit value identifying the mobile country code. One or two-digit entries will be prepended by 0 to create a three-digit value.
- Mobile Network Code—Defines the two or three-digit value identifying the network code.
- Add—Add the specified country code and network code to the IMSI Prefix table.
- Delete—Deletes the specified country code and network code from the IMSI Prefix table.

Add/Edit GTP Policy Map (Security Level)

Configuration > Global Objects > Inspect Maps > GTP > GTP Inspect Map > Basic View

The Add/Edit GTP Policy Map pane lets you configure the security level and additional settings for GTP application inspection maps.

Fields

- Name—When adding a GTP map, enter the name of the GTP map. When editing a GTP map, the name of the previously configured GTP map is shown.
- Description—Enter the description of the GTP map, up to 200 characters in length.
- Security Level—Security level low only.

    Do not Permit Errors

    Maximum Number of Tunnels: 500

    GSN timeout: 00:30:00

    Pdp-Context timeout: 00:30:00

    Request timeout: 00:01:00

    Signaling timeout: 00:30:00.

    Tunnel timeout: 01:00:00.

    T3-response timeout: 00:00:20.

    Drop and log unknown message IDs.

    – IMSI Prefix Filtering—Opens the IMSI Prefix Filtering dialog box to configure IMSI prefix filters.

    – Default Level—Sets the security level back to the default.

- Details—Shows the Parameters, IMSI Prefix Filtering, and Inspections tabs to configure additional settings.

Add/Edit GTP Policy Map (Details)

Configuration > Global Objects > Inspect Maps > GTP > GTP Inspect Map > Advanced View

The Add/Edit GTP Policy Map pane lets you configure the security level and additional settings for GTP application inspection maps.

Fields

- Name—When adding a GTP map, enter the name of the GTP map. When editing a GTP map, the name of the previously configured GTP map is shown.

- Description—Enter the description of the GTP map, up to 200 characters in length.

- Security Level—Shows the security level and IMSI prefix filtering settings to configure.

- Permit Parameters—Tab that lets you configure the permit parameters for the GTP inspect map.

    - Object Groups to Add

      From object group—Specify an object group or use the browse button to open the Add Network Object Group dialog box.

      To object group—Specify an object group or use the browse button to open the Add Network Object Group dialog box.

    - Add—Add the specified country code and network code to the IMSI Prefix table.

    - Delete—Deletes the specified country code and network code from the IMSI Prefix table.

    - Permit Errors—Lets any packets that are invalid or that encountered an error during inspection to be sent through the ASA instead of being dropped. By default, all invalid packets or packets that failed during parsing are dropped.

- General Parameters—Tab that lets you configure the general parameters for the GTP inspect map.

    - Maximum Number of Requests—Lets you change the default for the maximum request queue size allowed. The default for the maximum request queue size is 200. Specifies the maximum number of GTP requests that will be queued waiting for a response. The permitted range is from 1 to 9999999.

    - Maximum Number of Tunnels—Lets you change the default for the maximum number of tunnels allowed. The default tunnel limit is 500. Specifies the maximum number of tunnels allowed. The permitted range is from 1 to 9999999 for the global overall tunnel limit.

    - Timeouts

      GSN timeout—Lets you change the default for the maximum period of inactivity before a GSN is removed. The default is 30 minutes. Timeout is in the format *hh*:*mm*:*ss*, where *hh* specifies the hour, *mm* specifies the minutes, and *ss* specifies the seconds. A value 0 means never tear down.

      PDP-Context timeout—Lets you change the default for the maximum period of inactivity before receiving the PDP Context for a GTP session. The default is 30 minutes. Timeout is in the format *hh*:*mm*:*ss*, where *hh* specifies the hour, *mm* specifies the minutes, and *ss* specifies the seconds. A value 0 means never tear down.

      Request Queue—Lets you change the default for the maximum period of inactivity before receiving the GTP message during a GTP session. The default is 1 minute. Timeout is in the format *hh*:*mm*:*ss*, where *hh* specifies the hour, *mm* specifies the minutes, and *ss* specifies the seconds. A value 0 means never tear down.

Signaling—Lets you change the default for the maximum period of inactivity before a GTP signaling is removed. The default is 30 minutes. Timeout is in the format *hh*:*mm*:*ss*, where *hh* specifies the hour, *mm* specifies the minutes, and *ss* specifies the seconds. A value 0 means never tear down.

Tunnel—Lets you change the default for the maximum period of inactivity for the GTP tunnel. The default is 1 hour. Timeout is in the format *hh*:*mm*:*ss*, where *hh* specifies the hour, *mm* specifies the minutes, and *ss* specifies the seconds. A value 0 means never tear down Request timeout—Specifies the GTP Request idle timeout.

T3-Response timeout—Specifies the maximum wait time for a response before removing the connection.

- IMSI Prefix Filtering—Tab that lets you configure the IMSI prefix filtering for the GTP inspect map.
    - Mobile Country Code—Defines the non-zero, three-digit value identifying the mobile country code. One or two-digit entries will be prepended by 0 to create a three-digit value.
    - Mobile Network Code—Defines the two or three-digit value identifying the network code.
    - Add—Add the specified country code and network code to the IMSI Prefix table.
    - Delete—Deletes the specified country code and network code from the IMSI Prefix table.
- Inspections—Tab that lets you configure the GTP inspect maps.
    - Match Type—Shows the match type, which can be a positive or negative match.
    - Criterion—Shows the criterion of the GTP inspection.
    - Value—Shows the value to match in the GTP inspection.
    - Action—Shows the action if the match condition is met.
    - Log—Shows the log state.
    - Add—Opens the Add GTP Inspect dialog box to add an GTP inspection.
    - Edit—Opens the Edit GTP Inspect dialog box to edit an GTP inspection.
    - Delete—Deletes an GTP inspection.
    - Move Up—Moves an inspection up in the list.
    - Move Down—Moves an inspection down in the list.

Add/Edit GTP Map

Configuration > Global Objects > Inspect Maps > GTP > GTP Inspect Map > Add/Edit GTP Map

The Add/Edit GTP Inspect dialog box lets you define the match criterion and value for the GTP inspect map.

Fields

- Match Type—Specifies whether traffic should match or not match the values.

    For example, if No Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the class map.

- Criterion—Specifies which criterion of GTP traffic to match.
    - Access Point Name—Match on access point name.
    - Message ID—Match on the message ID.

- – Message Length—Match on the message length
- – Version—Match on the version.

- • Access Point Name Criterion Values—Specifies an access point name to be matched. By default, all messages with valid APNs are inspected, and any APN is allowed.
    - – Regular Expression—Lists the defined regular expressions to match.
    - – Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
    - – Regular Expression Class—Lists the defined regular expression classes to match.
    - – Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
    - – Action—Drop.
    - – Log—Enable or disable.

- • Message ID Criterion Values—Specifies the numeric identifier for the message that you want to match. The valid range is 1 to 255. By default, all valid message IDs are allowed.
    - – Value—Specifies whether value is an exact match or a range.

        Equals—Enter a value.

        Range—Enter a range of values.
    - – Action—Drop packet or limit rate (pps).
    - – Log—Enable or disable.

- • Message Length Criterion Values—Lets you change the default for the maximum message length for the UDP payload that is allowed.
    - – Minimum value—Specifies the minimum number of bytes in the UDP payload. The range is from 1 to 65536.
    - – Maximum value—Specifies the maximum number of bytes in the UDP payload. The range is from 1 to 65536.
    - – Action—Drop packet.
    - – Log—Enable or disable.

- • Version Criterion Values—Specifies the GTP version for messages that you want to match. The valid range is 0-255. Use 0 to identify Version 0 and 1 to identify Version 1. Version 0 of GTP uses port 3386, while Version 1 uses port 2123. By default all GTP versions are allowed.
    - – Value—Specifies whether value is an exact match or a range.

        Equals—Enter a value.

        Range—Enter a range of values.
    - – Action—Drop packet.
    - – Log—Enable or disable.

RADIUS Accounting Inspection

This section describes the RADIUS Accounting inspection engine. This section includes the following topics:

RADIUS Accounting Inspection Overview

One of the well known problems is the over-billing attack in GPRS networks. The over-billing attack can cause consumers anger and frustration by being billed for services that they have not used. In this case, a malicious attacker sets up a connection to a server and obtains an IP address from the SGSN. When the attacker ends the call, the malicious server will still send packets to it, which gets dropped by the GGSN, but the connection from the server remains active. The IP address assigned to the malicious attacker gets released and reassigned to a legitimate user who will then get billed for services that the attacker will use.

RADIUS accounting inspection prevents this type of attack by ensuring the traffic seen by the GGSN is legitimate. With the RADIUS accounting feature properly configured, the security appliance tears down a connection based on matching the Framed IP attribute in the Radius Accounting Request Start message with the Radius Accounting Request Stop message. When the Stop message is seen with the matching IP address in the Framed IP attribute, the security appliance looks for all connections with the source matching the IP address.

You have the option to configure a secret pre-shared key with the RADIUS server so the security appliance can validate the message. If the shared secret is not configured, the security appliance does not need to validate the source of the message and will only check that the source IP address is one of the configured addresses allowed to send the RADIUS messages.

Note   When using RADIUS accounting inspection with GPRS enabled, the ASA checks for the 3GPP-Session-Stop-Indicator in the Accounting Request STOP messages to properly handle secondary PDP contexts. Specifically, the ASA requires that the Accounting Request STOP messages include the 3GPP-SGSN-Address attribute before it will terminate the user sessions and all associated connections. Some third-party GGSNs might not send this attribute by default.

Select RADIUS Accounting Map

The Select RADIUS Accounting Map dialog box lets you select a defined RADIUS accounting map or define a new one.

Fields

- Add—Lets you add a new RADIUS accounting map.

Add RADIUS Accounting Policy Map

The Add RADIUS Accounting Policy Map dialog box lets you add the basic settings for the RADIUS accounting map.

Fields

- Name—Enter the name of the previously configured RADIUS accounting map.

- Description—Enter the description of the RADIUS accounting map, up to 100 characters in length.

- Host Parameters tab:

  – Host IP Address—Specify the IP address of the host that is sending the RADIUS messages.

  – Key: (optional)—Specify the key.

  – Add—Adds the host entry to the Host table.

  – Delete—Deletes the host entry from the Host table.

- Other Parameters tab:

  – Attribute Number—Specify the attribute number to validate when an Accounting Start is received.

  – Add—Adds the entry to the Attribute table.

  – Delete—Deletes the entry from the Attribute table.

  – Send response to the originator of the RADIUS message—Sends a message back to the host from which the RADIUS message was sent.

  – Enforce timeout—Enables the timeout for users.

  – Users Timeout—Timeout for the users in the database (hh:mm:ss).

RADIUS Inspect Map

The RADIUS pane lets you view previously configured RADIUS application inspection maps. A RADIUS map lets you change the default configuration values used for RADIUS application inspection. You can use a RADIUS map to protect against an overbilling attack.

Fields

- Name—Enter the name of the inspect map, up to 40 characters in length.

- Description—Enter the description of the inspect map, up to 200 characters in length.

- RADIUS Inspect Maps—Table that lists the defined RADIUS inspect maps. The defined inspect maps are also listed in the RADIUS area of the Inspect Maps tree.

- Add—Adds the new RADIUS inspect map to the defined list in the RADIUS Inspect Maps table and to the RADIUS area of the Inspect Maps tree. To configure the new RADIUS map, select the RADIUS entry in Inspect Maps tree.

- Delete—Deletes the application inspection map selected in the RADIUS Inspect Maps table and from the RADIUS area of the Inspect Maps tree.

RADIUS Inspect Map Host

The RADIUS Inspect Map Host Parameters pane lets you configure the host parameter settings for the inspect map.

Fields

- Name—Shows the name of the previously configured RADIUS accounting map.

- Description—Enter the description of the RADIUS accounting map, up to 200 characters in length.
- Host Parameters—Lets you configure host parameters.
  - Host IP Address—Specify the IP address of the host that is sending the RADIUS messages.
  - Key: (optional)—Specify the key.
- Add—Adds the host entry to the Host table.
- Delete—Deletes the host entry from the Host table.

RADIUS Inspect Map Other

The RADIUS Inspect Map Other Parameters pane lets you configure additional parameter settings for the inspect map.

Fields

- Name—Shows the name of the previously configured RADIUS accounting map.
- Description—Enter the description of the RADIUS accounting map, up to 200 characters in length.
- Other Parameters—Lets you configure additional parameters.
  - Send response to the originator of the RADIUS message—Sends a message back to the host from which the RADIUS message was sent.
  - Enforce timeout—Enables the timeout for users.

    Users Timeout—Timeout for the users in the database (hh:mm:ss).
  - Enable detection of GPRS accounting—Enables detection of GPRS accounting. This option is only available when GTP/GPRS license is enabled.
  - Validate Attribute—Attribute information.

    Attribute Number—Specify the attribute number to validate when an Accounting Start is received.

    Add—Adds the entry to the Attribute table.

    Delete—Deletes the entry from the Attribute table.

RSH Inspection

RSH inspection is enabled by default. The RSH protocol uses a TCP connection from the RSH client to the RSH server on TCP port 514. The client and server negotiate the TCP port number where the client listens for the STDERR output stream. RSH inspection supports NAT of the negotiated port number if necessary.

SNMP Inspection

This section describes the SNMP inspection engine. This section includes the following topics:

SNMP Inspection Overview

SNMP application inspection lets you restrict SNMP traffic to a specific version of SNMP. Earlier versions of SNMP are less secure; therefore, denying certain SNMP versions may be required by your security policy. The ASA can deny SNMP versions 1, 2, 2c, or 3. You control the versions permitted by creating an SNMP map.

You then apply the SNMP map when you enable SNMP inspection according to the Configuring Application Layer Protocol Inspection, page 8-7.

Select SNMP Map

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Select SNMP Map

The Select SNMP Map dialog box lets you select or create a new SNMP map. An SNMP map lets you change the configuration values used for SNMP application inspection. The Select SNMP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default SNMP inspection map—Specifies to use the default SNMP map.

- Select an SNMP map for fine control over inspection**—Lets you select a defined application inspection map or add a new one.

- Add—Opens the Add Policy Map dialog box for the inspection.

SNMP Inspect Map

Configuration > Global Objects > Inspect Maps > SNMP

The SNMP pane lets you view previously configured SNMP application inspection maps. An SNMP map lets you change the default configuration values used for SNMP application inspection.

Fields

- Map Name—Lists previously configured application inspection maps. Select a map and click **Edit** to view or change an existing map.

- Add—Configures a new SNMP inspect map.

- Edit—Edits the selected SNMP entry in the SNMP Inspect Maps table.

- Delete—Deletes the inspect map selected in the SNMP Inspect Maps table.

Add/Edit SNMP Map

Configuration > Global Objects > Inspect Maps > SNMP > Add/Edit SNMP Map (You can get to this dialog box through various paths.)

The Add/Edit SNMP Map dialog box lets you create a new SNMP map for controlling SNMP application inspection.

Fields

- SNMP Map Name—Defines the name of the application inspection map.

- SNMP version 1—Enables application inspection for SNMP version 1.
- SNMP version 2 (party based)—Enables application inspection for SNMP version 2.
- SNMP version 2c (community based)—Enables application inspection for SNMP version 2c.
- SNMP version 3—Enables application inspection for SNMP version 3.

XDMCP Inspection

XDMCP inspection is enabled by default; however, the XDMCP inspection engine is dependent upon proper configuration of the **established** command.

XDMCP is a protocol that uses UDP port 177 to negotiate X sessions, which use TCP when established.

For successful negotiation and start of an XWindows session, the ASA must allow the TCP back connection from the Xhosted computer. To permit the back connection, use the **established** command on the ASA. Once XDMCP negotiates the port to send the display, The **established** command is consulted to verify if this back connection should be permitted.

During the XWindows session, the manager talks to the display Xserver on the well-known port 6000 | n. Each display has a separate connection to the Xserver, as a result of the following terminal setting.

```
setenv DISPLAY Xserver:n
```

where *n* is the display number.

When XDMCP is used, the display is negotiated using IP addresses, which the ASA can NAT if needed. XDCMP inspection does not support PAT.

P A R T   4

Unified Communications

Information About the ASA in Cisco Unified Communications

This chapter describes how to configure the adaptive security appliance for Cisco Unified Communications Proxy features.

This chapter includes the following sections:

Information About the ASA in Cisco Unified Communications

This section describes the Cisco UC Proxy features. The purpose of a proxy is to terminate and reoriginate connections between a client and server. The proxy delivers a range of security functions such as traffic inspection, protocol conformance, and policy control to ensure security for the internal network. An increasingly popular function of a proxy is to terminate encrypted connections in order to apply security policies while maintaining confidentiality of connections. The ASA is a strategic platform to provide proxy functions for unified communications deployments.

The Cisco UC Proxy includes the following solutions:

Phone Proxy: Secure remote access for Cisco encrypted endpoints, and VLAN traversal for Cisco softphones

The phone proxy feature enables termination of Cisco SRTP/TLS-encrypted endpoints for secure remote access. The phone proxy allows large scale deployments of secure phones without a large scale VPN remote access hardware deployment. End-user infrastructure is limited to just the IP endpoint, without VPN tunnels or hardware.

The Cisco adaptive security appliance phone proxy is the replacement product for the Cisco Unified Phone Proxy. Additionally, the phone proxy can be deployed for voice/data VLAN traversal for softphone applications. Cisco IP Communicator (CIPC) traffic (both media and signaling) can be proxied through the ASA, thus traversing calls securely between voice and data VLANs.

For information about the differences between the TLS proxy and phone proxy, go to the following URL for Unified Communications content, including TLS Proxy vs. Phone Proxy white paper:

http://www.cisco.com/go/secureuc

TLS Proxy: Decryption and inspection of Cisco Unified Communications encrypted signaling

End-to-end encryption often leaves network security appliances "blind" to media and signaling traffic, which can compromise access control and threat prevention security functions. This lack of visibility can result in a lack of interoperability between the firewall functions and the encrypted voice, leaving businesses unable to satisfy both of their key security requirements.

The ASA is able to intercept and decrypt encrypted signaling from Cisco encrypted endpoints to the Cisco Unified Communications Manager (Cisco UCM), and apply the required threat protection and access control. It can also ensure confidentiality by re-encrypting the traffic onto the Cisco UCM servers.

Typically, the ASA TLS Proxy functionality is deployed in campus unified communications network. This solution is ideal for deployments that utilize end to end encryption and firewalls to protect Unified Communications Manager servers.

Mobility Proxy: Secure connectivity between Cisco Unified Mobility Advantage server and Cisco Unified Mobile Communicator clients

Cisco Unified Mobility solutions include the Cisco Unified Mobile Communicator (Cisco UMC), an easy-to-use software application for mobile handsets that extends enterprise communications applications and services to mobile phones and the Cisco Unified Mobility Advantage (Cisco UMA) server. The Cisco Unified Mobility solution streamlines the communication experience, enabling single number reach and integration of mobile endpoints into the Unified Communications infrastructure.

The security appliance acts as a proxy, terminating and reoriginating the TLS signaling between the Cisco UMC and Cisco UMA. As part of the proxy security functionality, inspection is enabled for the Cisco UMA Mobile Multiplexing Protocol (MMP), the protocol between Cisco UMC and Cisco UMA.

Presence Federation Proxy: Secure connectivity between Cisco Unified Presence servers and Cisco/Microsoft Presence servers

Cisco Unified Presence solution collects information about the availability and status of users, such as whether they are using communication devices, such as IP phones at particular times. It also collects information regarding their communications capabilities, such as whether web collaboration or video conferencing is enabled. Using user information captured by Cisco Unified Presence, applications such as Cisco Unified Personal Communicator and Cisco UCM can improve productivity by helping users connect with colleagues more efficiently through determining the most effective way for collaborative communication.

Using the ASA as a secure presence federation proxy, businesses can securely connect their Cisco Unified Presence (Cisco UP) servers to other Cisco or Microsoft Presence servers, enabling intra-enterprise communications. The security appliance terminates the TLS connectivity between the servers, and can inspect and apply policies for the SIP communications between the servers.

Cisco Intercompany Media Engine Proxy: Secure connectivity between Cisco UCM servers in different enterprises for IP Phone traffic

As more unified communications are deployed within enterprises, cases where business-to-business calls utilize unified communications on both sides with the Public Switched Network (PSTN) in the middle become increasingly common. All outside calls go over circuits to telephone providers and from there are delivered to all external destinations.

The Cisco Intercompany Media Engine gradually creates dynamic, encrypted VoIP connections between businesses, so that a collection of enterprises that work together end up looking like one giant business with secure VoIP interconnections between them.

There are three components to a Cisco Intercompany Media Engine deployment within an enterprise: a Cisco Intercompany Media Engine server, a call agent (the Cisco Unified Communications Manager) and an ASA running the Cisco Intercompany Media Engine Proxy.

The ASA provides perimeter security by encrypting signaling connections between enterprises and preventing unathorized calls. An ASA running the Cisco Intercompany Media Engine Proxy can either be deployed as an Internet firewall or be designated as a Cisco Intercompany Media Engine Proxy and placed in the DMZ, off the path of the regular Internet traffic.

TLS Proxy Applications in Cisco Unified Communications

Table 13-1 shows the Cisco Unified Communications applications that utilize the TLS proxy on the ASA.

Table 13-1   TLS Proxy Applications and the Security Appliance

Application	TLS Client	TLS Server	Client Authentication	Security Appliance Server Role	Security Appliance Client Role
Phone Proxy and TLS Proxy	IP phone	Cisco UCM	Yes	Proxy certificate, self-signed or by internal CA	Local dynamic certificate signed by the ASA CA (might not need certificate for phone proxy application)
Mobility Proxy	Cisco UMC	Cisco UMA	No	Using the Cisco UMA private key or certificate impersonation	Any static configured certificate
Presence Federation Proxy	Cisco UP or MS LCS/OCS	Cisco UP or MS LCS/OCS	Yes	Proxy certificate, self-signed or by internal CA	Using the Cisco UP private key or certificate impersonation

The ASA supports TLS proxy for various voice applications. For the phone proxy, the TLS proxy running on the ASA has the following key features:

- The ASA forces remote IP phones connecting to the phone proxy through the Internet to be in secured mode even when the Cisco UCM cluster is in non-secure mode.

- The TLS proxy is implemented on the ASA to intercept the TLS signaling from IP phones.

- The TLS proxy decrypts the packets, sends packets to the inspection engine for NAT rewrite and protocol conformance, optionally encrypts packets, and sends them to Cisco UCM or sends them in clear text if the IP phone is configured to be in nonsecure mode on the Cisco UCM.

- The ASA acts as a media terminator as needed and translates between SRTP and RTP media streams.

- The TLS proxy is a transparent proxy that works based on establishing trusted relationship between the TLS client, the proxy (the ASA), and the TLS server.

For the Cisco Unified Mobility solution, the TLS client is a Cisco UMA client and the TLS server is a Cisco UMA server. The ASA is between a Cisco UMA client and a Cisco UMA server. The mobility proxy (implemented as a TLS proxy) for Cisco Unified Mobility allows the use of an imported PKCS-12 certificate for server proxy during the handshake with the client. Cisco UMA clients are not required to present a certificate (no client authentication) during the handshake.

For the Cisco Unified Presence solution, the ASA acts as a TLS proxy between the Cisco UP server and the foreign server. This allows the ASA to proxy TLS messages on behalf of the server that initiates the TLS connection, and route the proxied TLS messages to the client. The ASA stores certificate trustpoints for the server and the client, and presents these certificates on establishment of the TLS session.

Licensing for Cisco Unified Communications Proxy Features

The Cisco Unified Communications proxy features supported by the ASA require a Unified Communications Proxy license:

- Phone proxy
- TLS proxy for encrypted voice inspection
- Presence federation proxy
- Intercompany media engine proxy

Note   In Version 8.2(2) and later, the Mobility Advantage proxy no longer requires a Unified Communications Proxy license.

The following table shows the Unified Communications Proxy license details by platform for the phone proxy, TLS proxy for encrypted voice inspection, and presence federation proxy:

Note   This feature is not available on No Payload Encryption models.

Model	License Requirement[1]
ASA 5505	Base License and Security Plus License: 2 sessions. *Optional license: 24 sessions.*
ASA 5512-X	Base License or Security Plus License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5515-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5525-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, or 1000 sessions.*
ASA 5545-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, or 2000 sessions.*
ASA 5555-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*

Model	License Requirement[1]
ASA 5585-X with SSP-10	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5585-X with SSP-20, -40, or -60	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*
ASASM	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*
ASAv with 1 Virtual CPU	Standard and Premium Licenses: 250 sessions.
ASAv with 4 Virtual CPUs	Standard and Premium Licenses: 1000 sessions.

1. The following applications use TLS proxy sessions for their connections. Each TLS proxy session used by these applications (and only these applications) is counted against the UC license limit:
- Phone Proxy
- Presence Federation Proxy
- Encrypted Voice Inspection

Other applications that use TLS proxy sessions do not count towards the UC limit, for example, Mobility Advantage Proxy (which does not require a license) and IME (which requires a separate IME license).

Some UC applications might use multiple sessions for a connection. For example, if you configure a phone with a primary and backup Cisco Unified Communications Manager, there are 2 TLS proxy connections, so 2 UC Proxy sessions are used.

You independently set the TLS proxy limit using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. When you apply a UC license that is higher than the default TLS proxy limit, the ASA automatically sets the TLS proxy limit to match the UC limit. The TLS proxy limit takes precedence over the UC license limit; if you set the TLS proxy limit to be less than the UC license, then you cannot use all of the sessions in your UC license.

Note: For license part numbers ending in "K8" (for example, licenses under 250 users), TLS proxy sessions are limited to 1000. For license part numbers ending in "K9" (for example, licenses 250 users or larger), the TLS proxy limit depends on the configuration, up to the model limit. K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted.

Note: If you clear the configuration, then the TLS proxy limit is set to the default for your model; if this default is lower than the UC license limit, then you see an error message to use the to raise the limit again (in ASDM, use the **TLS Proxy** pane). If you use failover and use **File > Save Running Configuration to Standby Unit** on the primary unit to force a configuration synchronization, the **clear configure all** command is generated on the secondary unit automatically, so you may see the warning message on the secondary unit. Because the configuration synchronization restores the TLS proxy limit set on the primary unit, you can ignore the warning.

You might also use SRTP encryption sessions for your connections:
- For K8 licenses, SRTP sessions are limited to 250.
- For K9 licenses, there is not limit.

Note: Only calls that require encryption/decryption for media are counted towards the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count towards the limit.

Table 13-2 shows the default and maximum TLS session details by platform.

Table 13-2   Default and Maximum TLS Sessions on the Security Appliance

Security Appliance Platform	Default TLS Sessions	Maximum TLS Sessions
ASA 5505	10	80

The following table shows the Unified Communications Proxy license details by platform for intercompany media engine proxy:

Note   This feature is not available on No Payload Encryption models.

Model	License Requirement
All models	Intercompany Media Engine license.
	When you enable the Intercompany Media Engine (IME) license, you can use TLS proxy sessions up to the configured TLS proxy limit. If you also have a Unified Communications (UC) license installed that is higher than the default TLS proxy limit, then the ASA sets the limit to be the UC license limit plus an additional number of sessions depending on your model. You can manually configure the TLS proxy limit using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. If you also install the UC license, then the TLS proxy sessions available for UC are also available for IME sessions. For example, if the configured limit is 1000 TLS proxy sessions, and you purchase a 750-session UC license, then the first 250 IME sessions do not affect the sessions available for UC. If you need more than 250 sessions for IME, then the remaining 750 sessions of the platform limit are used on a first-come, first-served basis by UC and IME.
	• For a license part number ending in "K8", TLS proxy sessions are limited to 1000.
	• For a license part number ending in "K9", the TLS proxy limit depends on your configuration and the platform model.
	Note K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted.
	You might also use SRTP encryption sessions for your connections:
	• For a K8 license, SRTP sessions are limited to 250.
	• For a K9 license, there is no limit.
	Note Only calls that require encryption/decryption for media are counted toward the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count toward the limit.

For more information about licensing, see the general operations configuration guide.

CHAPTER **14**

Cisco Unified Communication Wizard

This chapter describes how to configure the ASA for Cisco Unified Communications Proxy features.

This chapter includes the following sections:

Information about the Cisco Unified Communication Wizard

> **Note** The Unified Communication Wizard is supported for the ASA version 8.3(1) and later.

The Unified Communication Wizard assists you in configuring the following Unified Communications proxies on the ASA:

- Cisco Phone Proxy

  See Configuring the Phone Proxy by using the Unified Communication Wizard, page 14-4.
- Cisco Mobility Advantage Proxy

  See Configuring the Mobility Advantage by using the Unified Communication Wizard, page 14-11.
- Cisco Presence Federation Proxy

  See Configuring the Presence Federation Proxy by using the Unified Communication Wizard, page 14-14.
- Cisco Intercompany Media Engine Proxy

  See Configuring the UC-IME by using the Unified Communication Wizard, page 14-16.

The wizard simplifies the configuration of the Unified Communications proxies in the following ways:

- You enter all required data in the wizard steps. You are not required to navigate various ASDM screens to configure the Unified Communications proxies.

- The wizard generates configuration settings for the Unified Communications proxies where possible, automatically, without requiring you to enter data. For example, the wizard configures the required ACLs, IP address translation (NAT and PAT) statements, self-signed certificates, TLS proxies, and application inspection.

- The wizard displays network diagrams to illustrate data collection.

To access the Unified Communication Wizard, choose one of the following paths in the main ASDM application window:

- **Wizards > Unified Communication Wizard**.

- **Configuration > Firewall > Unified Communications,** and then click **Unified Communication Wizard**.

Phone Proxy: Secure remote access for Cisco encrypted endpoints, and VLAN traversal for Cisco softphones

The phone proxy feature enables termination of Cisco SRTP/TLS-encrypted endpoints for secure remote access. The phone proxy allows large scale deployments of secure phones without a large scale VPN remote access hardware deployment. End-user infrastructure is limited to just the IP endpoint, without VPN tunnels or hardware.

The Cisco adaptive security appliance phone proxy is the replacement product for the Cisco Unified Phone Proxy. Additionally, the phone proxy can be deployed for voice/data VLAN traversal for softphone applications. Cisco IP Communicator (CIPC) traffic (both media and signaling) can be proxied through the ASA, thus traversing calls securely between voice and data VLANs.

For information about the differences between the TLS proxy and phone proxy, go to the following URL for Unified Communications content, including TLS Proxy vs. Phone Proxy white paper:

http://www.cisco.com/go/secureuc

Mobility Advantage Proxy: Secure connectivity between Cisco Mobility Advantage server and Cisco Unified Mobile Communicator clients

Cisco Mobility Advantage solutions include the Cisco Unified Mobile Communicator (Cisco UMC), an easy-to-use software application for mobile handsets that extends enterprise communications applications and services to mobile phones and the Cisco Unified Mobility Advantage (Cisco UMA) server. The Cisco Mobility Advantage solution streamlines the communication experience, enabling single number reach and integration of mobile endpoints into the Unified Communications infrastructure.

The security appliance acts as a proxy, terminating and reoriginating the TLS signaling between the Cisco UMC and Cisco UMA. As part of the proxy security functionality, inspection is enabled for the Cisco UMA Mobile Multiplexing Protocol (MMP), the protocol between Cisco UMC and Cisco UMA.

Presence Federation Proxy: Secure connectivity between Cisco Unified Presence servers and Cisco/Microsoft Presence servers

Cisco Unified Presence solution collects information about the availability and status of users, such as whether they are using communication devices, such as IP phones at particular times. It also collects information regarding their communications capabilities, such as whether web collaboration or video conferencing is enabled. Using user information captured by Cisco Unified Presence, applications such as Cisco Unified Personal Communicator and Cisco UCM can improve productivity by helping users connect with colleagues more efficiently through determining the most effective way for collaborative communication.

Using the ASA as a secure presence federation proxy, businesses can securely connect their Cisco Unified Presence (Cisco UP) servers to other Cisco or Microsoft Presence servers, enabling intra-enterprise communications. The security appliance terminates the TLS connectivity between the servers, and can inspect and apply policies for the SIP communications between the servers.

Cisco Intercompany Media Engine Proxy: Secure connectivity between Cisco UCM servers in different enterprises for IP Phone traffic

As more unified communications are deployed within enterprises, cases where business-to-business calls utilize unified communications on both sides with the Public Switched Network (PSTN) in the middle become increasingly common. All outside calls go over circuits to telephone providers and from there are delivered to all external destinations.

The Cisco Intercompany Media Engine (UC-IME) gradually creates dynamic, encrypted VoIP connections between businesses, so that a collection of enterprises that work together end up looking like one giant business with secure VoIP interconnections between them.

There are three components to a Cisco Intercompany Media Engine deployment within an enterprise: a Cisco Intercompany Media Engine server, a call agent (the Cisco Unified Communications Manager) and an ASA running the Cisco Intercompany Media Engine Proxy.

The ASA provides perimeter security by encrypting signaling connections between enterprises and preventing unauthorized calls. An ASA running the Cisco Intercompany Media Engine Proxy can either be deployed as an Internet firewall or be designated as a Cisco Intercompany Media Engine Proxy and placed in the DMZ, off the path of the regular Internet traffic.

Licensing Requirements for the Unified Communication Wizard

To run the Unified Communication Wizard in ASDM, you require the following license:

Model	License Requirement
All models	Base License

However, to run each of the Unified Communications proxy features created by the wizard, you must have the appropriate Unified Communications Proxy licenses.

The Cisco Unified Communications proxy features supported by the ASA require a Unified Communications Proxy license:

- Cisco Phone Proxy
- TLS proxy for encrypted voice inspection
- Presence Federation Proxy
- Cisco Intercompany Media Engine Proxy

See Licensing for Cisco Unified Communications Proxy Features, page 13-4 for more information.

Note   The Cisco Intercompany Media Engine Proxy does not appear as an option in the Unified Communication Wizard unless the license required for this proxy is installed on the ASA.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

IPv6 Guidelines

Supports IPv6 addresses.

Additional Guidelines and Limitations

Using the Unified Communication Wizard to create the Unified Communications proxies has the following limitations and requirements:

- You must configure at least two interfaces on the ASA to use the UC Wizard to configure a Unified Communications proxy.

- For all Unified Communications proxies to function correctly, you must synchronize the clock on the ASA and all servers associated with each proxy, such as the Cisco Unified Communication Manager server, the Cisco Mobility Advantage server, the Cisco Unified Presence server, and the Cisco Intercompany Media Engine server.

- When you configure the Cisco Intercompany Media Engine Proxy for an off-path deployment, you must ensure that the public IP addresses and ports of the Cisco Unified Communications Manager servers and the public IP address for the media termination address are accessible from the Internet. The summary page of the Unified Communication Wizard reminds you of the requirements.

- If the ASA on which you configure the Cisco Mobility Advantage Proxy and the Cisco Presence Federation Proxy is located behind another firewall, you must ensure that the public IP addresses for the Cisco Mobility Advantage server and the Cisco Unified Presence server are accessible from the Internet.

- If you use the Unified Communication Wizard to create to the Presence Federation Proxy and the Cisco Intercompany Media Engine Proxy, you might be required to adjust the configuration of the ACLs created automatically by the wizard for each proxy. See Chapter 18, "ASA and Cisco Unified Presence" and Chapter 19, "ASA and Cisco Intercompany Media Engine Proxy", respectively, for information about the ACL requirements required by each proxy.

Configuring the Phone Proxy by using the Unified Communication Wizard

To configure the Cisco Unified Presence proxy by using ASDM, choose Wizards > Unified Communications Wizard from the menu. The Unified Communications Wizard opens. From the first page, select the Phone Proxy option under the Remote Access section.

The wizard automatically creates the necessary TLS proxy, then guides you through creating the Phone Proxy instance, importing and installing the required certificates, and finally enables the SIP and SCCP inspection for the Phone Proxy traffic automatically.

> **Note**   Any configuration created by the wizard should be maintained through the wizard to ensure proper synchronization. For example, if you create a phone proxy configuration through the UC wizard and then modify the configuration outside of the wizard, the rest of the wizard configuration is not updated, and the wizard configuration is not synchronized.
>
> Therefore, if you choose to change some part of the phone proxy configuration outside of the wizard, it is your responsibility to keep the rest of the configuration in synchronization.

The wizard guides you through four steps to configure the Phone Proxy:

Step 1   Select the Phone Proxy option.

Step 2   Specify settings to define the Cisco Unified Communications Manager (UCM) servers and TFTP servers, such the IP address and the address translation settings of each server, and the Cisco UCM cluster security mode. See Configuring the Private Network for the Phone Proxy, page 14-5 and Configuring Servers for the Phone Proxy, page 14-6.

Step 3   If required, enable Certificate Authority Proxy Function (CAPF). See Enabling Certificate Authority Proxy Function (CAPF) for IP Phones, page 14-8.

Step 4   Configure the public IP phone network, such as address translation settings for remote IP phones, whether to enable service setting for IP phones, and the HTTP proxy used by the IP phones. Configuring the Public IP Phone Network, page 14-9

Step 5   Specify the media termination address settings of the Cisco UCM. Configuring the Media Termination Address for Unified Communication Proxies, page 14-10.

The wizard completes by displaying a summary of the configuration created for Phone Proxy.

Configuring the Private Network for the Phone Proxy

The values that you specify in this page configure the connection from the ASA to the Cisco UCMs and TFTP servers by creating the necessary address translation settings and access control list entries.

Additionally, you specify the security mode for the Cisco UCM cluster. In a nonsecure cluster mode or a mixed mode where the phones are configured as nonsecure, the phone proxy behaves in the following ways:

- The TLS connections from the phones are terminated on the ASA and a TCP connection is initiated to the Cisco UCM.
- SRTP sent from external IP phones to the internal network IP phone via the ASA is converted to RTP.

In a mixed mode cluster where the internal IP phones are configured as authenticated, the TLS connection is not converted to TCP to the Cisco UCM but the SRTP is converted to RTP.

In a mixed mode cluster where the internal IP phone is configured as encrypted, the TLS connection remains a TLS connection to the Cisco UCM and the SRTP from the remote phone remains SRTP to the internal IP phone.

Step 1   From the Interface drop-down list, choose the interface on which the ASA listens for the Cisco UCM servers and TFTP servers. The Cisco UCM servers and TFTP servers must reside on the same interface.

Step 2   Specify each entity in the network (all Cisco UCM and TFTP servers) that the IP phones must trust. Click **Add** to add the servers. See Configuring Servers for the Phone Proxy, page 14-6.

To modify the configuration of a server already added to the configuration, select the server in the table and click **Edit**. The Edit Server dialog appears. See Configuring Servers for the Phone Proxy, page 14-6. At least one Cisco UCM and at least one TFTP server must be configured for the phone proxy.

Step 3   Specify the security mode of the Cisco UCM cluster by clicking one of the following options in the Unified CM Cluster Mode field:

- Non-secure—Specifies the cluster to be in nonsecure mode when configuring the Phone Proxy feature.

- Mixed—Specifies the cluster to be in mixed mode when configuring the Phone Proxy feature.

    If you selected the Mixed security mode, the Generate and Export LDC Certificate button becomes available.

Step 4   For a Mixed security mode only, configure local dynamic certificates (LDC) for the IP phones by performing the following steps:

- **a.** Click the **Generate and Export LDC Certificate** button.

    A dialog box appears stating "Enrollment succeeded," which indicates that the LDC was generated.

- **b.** Click **OK** to close the Enrollment Status dialog box. The Export certificate dialog box appears.

- **c.** In the Export to File field, enter the file name and path for the LDC or click browse to locate and select an existing file.

- **d.** Click the **Export Certificate** button. A dialog box appears indicating that the file was exported successfully.

- **e.** Click **OK** to close the dialog box. A dialog box appears reminding you to install the LDC on the Cisco UCMs.

- **f.** Click **OK** to close the dialog box.

    Once configured, the ASA presents this unique, dynamically-created certificate to the Cisco UCM on behalf of the IP phones.

Step 5   Click **Next**.

Configuring Servers for the Phone Proxy

The values that you specify in this page generate address translation settings, access list entries, trustpoints, and the corresponding CTL file entries for each server.

You must add a server for each entity in the network that the IP phones must trust. These servers include all Cisco UCM servers in the cluster and all the TFTP servers.

You must add at least one TFTP server and at least one Cisco UCM server for the phone proxy. You can configure up to five TFTP servers for the phone proxy. The TFTP server is assumed to be behind the firewall on the trusted network; therefore, the phone proxy intercepts the requests between the IP phones and TFTP server.

Note   When you delete a TFTP server from the Server list in Step 2 of the wizard, ASDM deletes only the TFTP server IP address from the configuration and does not remove from the configuration all the ACLs, NAT statements, object groups, etc. attached to the TFTP server. To remove those attached configuration

statements, you must delete them manually by using the appropriate area of ASDM or rerun the Unified Communications wizard without making any changes and apply the configuration to to remove these statements.

The servers that the IP phones must trust can be deployed on the network in one of the following ways:

- All the services required by the Cisco UCM server, namely the Cisco UCM, TFTP, and CAPF services, are running on one server. In this deployment, only one instance of each service exists. For this deployment, you can select Unified CM+ TFTP as the server type. You can either use Address only or Address and ports for address translation. Cisco recommends that you specify Address and ports for increased security.

- Deployments for larger enterprises might have redundant Cisco UCMs and dedicated servers for TFTP and CAPF services. In that type of deployment, use Address only for voice address translation and Address only or Address and ports for TFTP.

Table 14-1 lists the ports that are configured for Address and port translation by default:

Table 14-1   Port Configuration

Address	Default Port	Description
TFTP Server	69	Allows incoming TFTP
Cisco UCM	2000	Allows incoming non-secure SCCP
Cisco UCM	2443	Allows incoming secure SCCP
Cisco UCM	5061	Allows incoming secure SIP

Step 1 In the Server Type field, select the server from the drop-down list: Unified CM, TFTP, or Unified CM + TFTP. Select Unified CM + TFTP when the Cisco UCM and TFTP server reside on the same device.

Note Depending on which type of server you select (Unified CM or TFTP), only the necessary fields in this dialog box become available. Specifically, if the server type is Unified CM, the TFTP section in the dialog is unavailable. If the server type is TFTP, the Voice section is unavailable.

Step 2 In the Private Address field, specify the actual internal IP address of the server.

Step 3 In the FQDN field, enter the fully-qualified domain name of the server, which includes the hostname and domain name; for example, `ucm.cisco.com` (where `ucm` is the hostname and `cisco.com` is the domain name).

If you are configuring a Unified CM server, enter the fully-qualified domain name configured on the Cisco UCM.

If you are configuring a TFTP server, only specify the TFTP server fully-qualified domain name when that server is configured with FQDN. If the TFTP server is not configured with FQDN, you can leave the field blank.

Note Entering the fully-qualified domain name allows the ASA to perform hostname resolution when DNS lookup is not configured on the ASA or the configured DNS servers are unavailable.See the command reference for information about the **dns domain-lookup** command.

Step 4 In the Address Translation section, select whether to use the interface IP address or to enter a different IP address.

Selecting the Use interface IP radio button configures the server to use the IP address of the public interface. You select the public interface in step 4 of the wizard when you configure the public network for the phone proxy.

If the Use interface IP radio button is selected, you must specify port translation settings in the Voice and TFTP sections. Address-only translation is available only when you specify an IP address other than the IP address of the public interface.

When you select the Address only radio button, the ASA performs address translation on all traffic between the server and the IP phones. Selecting the Address and ports radio button limits address translation to the specified ports.

Step 5   (Unified CM or Unified CM + TFTP servers only) In the Voice section, configure inspection of SIP or SCCP protocol traffic, or both SIP and SCCP protocol traffic by completing the following fields:

    a.   In the Translation Type field, specify whether to use the Address only or the Address and ports.

    When the deployment has redundant Cisco UCM servers and dedicated servers for TFTP and CAPF services, select Address only for voice address translation.

    Select the Address and ports option when you want to limit address translation to the specified ports.

    b.   In the Voice Protocols field, select the inspection protocols supported by the IP phones deployed in the enterprise. Depending on which inspection protocols you select—SCCP, SIP, or SCCP and SIP—only the ports fields for the selected voice protocols are available.

    c.   In the Port Translation section, enter the private and public ports for the voice protocols.

    The default values for the voice ports appear in the text fields. If necessary, change the private ports to match the settings on the Cisco UCM. The values you set for the public ports are used by the IP phones to traverse the ASA and communicate with the Cisco UCM.

    The secure SCCP private port and public port are automatically configured. These port numbers are automatically set to the value of the non-secure port number plus 443.

Step 6   (TFTP or Unified CM + TFTP servers only) In the TFTP section, you can select either Address only or Address and port for address translation. Cisco recommends that you specify Address and port for increased security. Specifying Address and port configures the TFTP server to listen on port 69 for TFTP requests.

When the server type is Unified CM + TFTP, the wizard configures the same type of address translation for Voice and TFTP; for example, when the server type is Unified CM + TFTP and the Address only option is selected, the wizard creates a global address translation rule for all traffic to and from the server. In this case, configuring port translation for the TFTP server would be redundant.

Step 7   Click **OK** to add the server to the phone proxy configuration and return to step 2 of the wizard.

Enabling Certificate Authority Proxy Function (CAPF) for IP Phones

As an alternative to authenticating remote IP phones through the TLS handshake, you can configure authentication via locally significant certificate (LSC) provisioning. With LSC provisioning, you create a password for each remote IP phone user and each user enters the password on the remote IP phones to retrieve the LSC.

Because using LSC provisioning to authenticate remote IP phones requires the IP phones first register in nonsecure mode, Cisco recommends LSC provisioning be done inside the corporate network before giving the IP phones to end-users. Otherwise, having the IP phones register in nonsecure mode requires the Administrator to open the nonsecure signaling port for SIP and SCCP on the ASA.

See also the Cisco Unified Communications Manager Security Guide for information on Using the Certificate Authority Proxy Function (CAPF) to install a locally significant certificate (LSC).

If your network includes Cisco IP Communicators (CIPC) or you have LSC enabled IP phones, you must import the CAPF certificate from the Cisco UCM. The certificate will be used to generate the LSC on the IP phones.

If the Cisco UCM has more than one CAPF certificate, you must import all of them to the ASA. However, the wizard supports configuring only one CAPF certificate, which is the default. To import more than one CAPF certificate, go to Configuration > Device Management > Certificate Management > Identity Certificates.

You can configure LSC provisioning for additional end-user authentication. See the Cisco Unified Communications Manager configuration guide for information.

Step 1   Check the **Enable Certificate Authority Proxy Function** check box. The remaining fields in the page become available.

Step 2   Enter the private IP address of the LSC provider.

Step 3   In the Public Address field, specify whether to use the IP address of the ASA public interface or enter an IP address.

Specifying the private and public IP addresses for the LSC provider, creates an access list entry that allows the IP phones to contact the Cisco UCM by opening the CAPF port for LSC provisioning.

Step 4   In the Translation Type field, select the Address only or Address and ports radio button.

The IP phones must contact the CAPF service on the Cisco UCM. The address translation type (Address only versus Address and ports) you select for CAPF must match the address translation type of the Cisco UCM on which the CAPF service is running. You set the address translation type for that Cisco UCM server in the previous step of this wizard (see Configuring Servers for the Phone Proxy, page 14-6),

By default, the CAPF Service uses port 3804. Modify this default value only when it is modified on the Cisco UCM.

Step 5   If you selected the Address and ports radio button, enter the private and public ports for the CAPF service.

Step 6   Click the **Install CAPF Certificate** button. The Install Certificate dialog box appears. See Installing a Certificate, page 14-23.

Step 7   Click **Next**.

Configuring the Public IP Phone Network

The values that you specify in this page generate the address translation rules used for the IP phones and configure how the ASA handles IP phone settings.

Step 1   From the Interface drop-down list, choose the interface on which the ASA listens for connections from IP phones.

Step 2   To preserve Call Manager configuration on the IP phones, check the Preserve the Unified CM's configuration on the phone's service check box. When this check box is uncheck, the following service settings are disabled on the IP phones:

- Web Access

- PC Port

- Voice VLAN access

- Gratuitous ARP

- Span to PC Port

Step 3   To configure address translation for IP phones, check the Enable address translation for IP phones check box. Select whether to use the IP address of the ASA private interface (which you selected in step 2 of the wizard) or enter an IP address.

Configuring address translation for IP phone configures the address used by the IP phones. All traffic from the outside network converges into one source IP address so that, if there is another corporate firewall in the network, a pinhole needs to be opened only for that IP address rather than for all traffic.

Step 4   To configure an HTTP proxy for the Phone Proxy feature that is written into the IP phone's configuration file under the <proxyServerURL> tag, do the following:

a.   Check the Configure an HTTP proxy to redirect phone URLs... check box.

b.   In the IP Address field, type the IP address of the HTTP proxy

c.   In the Port field, enter the listening port of the HTTP proxy.

The IP address you enter should be the global IP address based on where the IP phone and HTTP proxy server is located. You can enter a hostname in the IP Address field when that hostname can be resolved to an IP address by the adaptive security appliance (for example, DNS lookup is configured) because the adaptive security appliance will resolve the hostname to an IP address. If a port is not specified, the default will be 8080.

d.   In the Interface field, select the interface on which the HTTP proxy resides on the adaptive security appliance.

Setting the proxy server configuration option for the Phone Proxy allows for an HTTP proxy on the DMZ or external network in which all the IP phone URLs are directed to the proxy server for services on the phones. This setting accommodates nonsecure HTTP traffic, which is not allowed back into the corporate network.

Step 5   Click **Next**.

Configuring the Media Termination Address for Unified Communication Proxies

The data from this step generates the MTA instance to be added to the Phone Proxy and the UC-IME proxy.

The phone proxy and the UC-IME proxy use the media termination address for Secure RTP (SRTP) and RTP traffic. SRTP traffic sent from external IP phones to the internal network IP phone via the ASA is converted to RTP traffic. The traffic is terminated on the adaptive security appliance. SRTP provides message authentication and replay protection to Internet media traffic such as audio and video. RTP defines a standardized packet format for delivering audio and video over the Internet.

For the UC-IME proxy and the Phone Proxy to be fully functional, you must ensure that the public IP address for the media termination address (MTA) is accessible from the Internet. The summary page of the Unified Communication Wizard reminds you of this requirement.

The MTA IP addresses that you specify must meet specific requirements. See Media Termination Instance Prerequisites, page 15-6 for information.

Step 1   In the field for the private IP address, enter the IP address on which private media traffic terminates. The IP address must be within the same subnet as the private interface IP address. The correct subnet range is provided to the right of the field for the private IP address.

Step 2   In the field for the public IP address, enter the IP address on which public media traffic terminates. The IP address must be within the same subnet as the public interface IP address. The correct subnet range is provided to the right of the field for the public IP address.

Step 3   Specify the minimum and maximum values for the RTP port range for the media termination instance.

Port values must be within the range of 1024 to 65535.

Step 4   Click **Next**.

The wizard completes by displaying a summary of the configuration created for proxy.

Configuring the Mobility Advantage by using the Unified Communication Wizard

Note   The Unified Communication Wizard is supported for the ASA version 8.3(1) and later.

The Unified Communication wizard guides you through the steps to configure the Mobility Advantage proxy. Choose **Wizards > Unified Communication Wizard** from the menu. The Unified Communication Wizard opens. Click the Cisco Mobility Advantage Proxy radio button under the Remote Access section.

When using the wizard to create the Mobility Advantage proxy, ASDM automatically creates the necessary TLS proxies, enables MMP inspection for the Mobility Advantage traffic, generates address translation (NAT) statements, and creates the access rules that are necessary to allow traffic between the Cisco Mobility Advantage server and the mobility clients.

The following steps provide the high-level overview for configuring the Mobility Advantage proxy:

Step 1   Specify settings to define the private and public network topology, such the public and private network interfaces, and the IP addresses of the Cisco Mobility Advantage server. See Configuring the Topology for the Cisco Mobility Advantage Proxy, page 14-12.

Step 2   Configure the certificates that are exchanged between the Cisco Mobility Advantage server and the ASA. See Configuring the Server-Side Certificates for the Cisco Mobility Advantage Proxy, page 14-12.

Step 3   Configure the client-side certificate management, namely the certificates that are exchanged between the Unified Mobile Communicator clients and the ASA. See Configuring the Client-Side Certificates for the Cisco Mobility Advantage Proxy, page 14-13.

The wizard completes by displaying a summary of the configuration created for Mobility Advantage Proxy.

Configuring the Topology for the Cisco Mobility Advantage Proxy

When configuring the Mobility Advantage Proxy, you specify settings to define the private and public network topology, such the private and public network interfaces, and the private and public IP addresses of the Cisco Mobility Advantage server.

The values that you specify in this page generate the following configuration settings for the Mobility Advantage Proxy:

- Static PAT for the Cisco Mobility Advantage server

- Static NAT for Cisco Unified Mobile Communicator clients if the Enable address translation for Mobility clients check box is checked.

- ACLs to allow Cisco Unified Mobile Communicator clients to access the Cisco Mobility Advantage server

Step 1 In the Private Network area, choose the interface from the drop-down list.

Step 2 In the Unified MA Server area, enter the private and public IP address for the Cisco Mobility Advantage server. Entering ports for these IP addresses is optional. By default port number 5443 is entered, which is the default TCP port for MMP inspection.

Step 3 In the FQDN field, enter the domain name for the Cisco Mobility Advantage server. This domain name is included in the certificate signing request that you generate later in this wizard.

Step 4 In the Public Network area, choose an interface from the drop-down list.

The proxy uses this interface for configuring static PAT for the Cisco Mobility Advantage server and the ACLs to allow Cisco Unified Mobile Communicator clients to access the Cisco Mobility Advantage server.

Step 5 **To configure whether address translation (NAT) is used by** Cisco Unified Mobile Communicator clients, check the **Enable address translation for Mobility clients** check box and choose whether to use the IP address of the public interface or whether to enter an IP address.

Step 6 Click **Next**.

Configuring the Server-Side Certificates for the Cisco Mobility Advantage Proxy

A trusted relationship between the ASA and the Cisco UMA server can be established with self-signed certificates. The ASA's identity certificate is exported, and then uploaded on the Cisco UMA server truststore. The Cisco UMA server certificate is downloaded, and then uploaded on the ASA truststore.

The supports using self-signed certificates only at this step.

Step 1 In the ASA's Identity Certificate area, click **Generate and Export ASA's Identity Certificate**.

An information dialog boxes appear indicating that the enrollment seceded. In the Enrollment Status dialog box, click **OK**. The Export certificate dialog box appears.

✎
Note • If an identity certificate for the ASA has already been created, the button in this area appears as **Export ASA's Identity Certificate** and the Export certificate dialog box immediately appears.

- When using the wizard to configure the Cisco Mobility Advantage proxy, the wizard only supports installing self-signed certificates.

Step 2   Export the identity certificate generated by the wizard for the ASA. See Exporting an Identity Certificate, page 14-23.

Step 3   In the Unified MA Server's Certificate area, click **Install Unified MA Server's Certificate**. The Install Certificate dialog appears.

Step 4   Locate the file containing the Cisco Mobility Advantage server certificate or paste the certificate details in the dialog box. See Installing a Certificate, page 14-23.

Step 5   Click **Next**.

> ✎
> **Note**   See the Cisco Mobility Advantage server documentation for information on how to export the certificate for this server.

Configuring the Client-Side Certificates for the Cisco Mobility Advantage Proxy

To establish a trust relationship between the Cisco Unified Mobile Communicator (UMC) clients and the ASA, the ASA uses a CA-signed certificate that is configured with the Cisco Mobility Advantage server's FQDN (also referred to as certificate impersonation).

In the Client-Side Certificate Management page, you enter both the intermediate CA certificate (if applicable, as in the cases of Verisign) and the signed ASA identity certificate.

> ✎
> **Note**   If the ASA already has a signed identity certificate, you can skip Step 1 in this procedure and proceed directly to Step 2.

Step 1   In the ASA's Identity Certificate area, click **Generate CSR**. The CSR parameters dialog box appears.

For information about specifying additional parameters for the certificate signing request (CSR), see Generating a Certificate Signing Request (CSR) for a Unified Communications Proxy, page 14-24.

Information dialog boxes appear indicating that the wizard is delivering the settings to the ASA and retrieving the certificate key pair information. The Identity Certificate Request dialog box appears.

For information about saving the CSR that was generated and submitting it to a CA, see Saving the Identity Certificate Request, page 14-25.

Step 2   Click **Install ASA's Identity Certificate**. Install the certificate. See Installing the ASA Identity Certificate on the Mobility Advantage Server, page 14-26.

Step 3   Click **Install Root CA's Certificate**. The Install Certificate dialog box appears. Install the certificate. See Installing a Certificate, page 14-23.

Step 4   Click **Next**.

The wizard completes by displaying a summary of the configuration created for Mobility Advantage Proxy.

Configuring the Presence Federation Proxy by using the Unified Communication Wizard

> **Note** The Unified Communication Wizard is supported for the ASA version 8.3(1) and later.

To configure the Cisco Unified Presence proxy by using ASDM, choose **Wizards > Unified Communication Wizard** from the menu. The Unified Communication Wizard opens. From the first page, select the Cisco Unified Presence Proxy option under the Business-to-Business section.

When using the wizard to create the Cisco Presence Federation proxy, ASDM automatically creates the necessary TLS proxies, enables SIP inspection for the Presence Federation traffic, generates address translation (static PAT) statements for the local Cisco Unified Presence server, and creates ACLs to allow traffic between the local Cisco Unified Presence server and remote servers.

The following steps provide the high-level overview for configuring the Presence Federation Proxy:

Step 1 Specify settings to define the private and public network topology, such the private and public IP address of the Presence Federation server. See Configuring the Topology for the Cisco Presence Federation Proxy, page 14-14.

Step 2 Configure the local-side certificate management, namely the certificates that are exchanged between the local Unified Presence Federation server and the ASA. See Configuring the Local-Side Certificates for the Cisco Presence Federation Proxy, page 14-15.

Step 3 Configure the remote-side certificate management, namely the certificates that are exchanged between the remote server and the ASA. See Configuring the Remote-Side Certificates for the Cisco Presence Federation Proxy, page 14-15.

The wizard completes by displaying a summary of the configuration created for the Presence Federation proxy.

Configuring the Topology for the Cisco Presence Federation Proxy

When configuring the Presence Federation Proxy, you specify settings to define the private and public network topology, such the private and public network interfaces, and the private and public IP addresses of the Cisco Unified Presence server.

The values that you specify in this page generate the following configuration settings for the Presence Federation Proxy:

- Static PAT for the local Cisco Unified Presence server
- ACLs for traffic between the local Cisco Unified Presence server and remote servers

Step 1 In the Private Network area, choose the interface from the drop-down list.

Step 2 In the Unified Presence Server area, enter the private and public IP address for the Unified Presence server. Entering ports for these IP addresses is optional. By default port number 5061 is entered, which is the default TCP port for SIP inspection.

Step 3   In the FQDN field, enter the domain name for the Unified Presence server. This domain name is included in the certificate signing request that you generate later in this wizard.

Step 4   In the Public Network area, choose the interface of the public network from the drop-down list. The proxy uses this interface for configuring static PAT for the local Cisco Unified Presence server and for configuring ACLs to allow remote servers to access the Cisco Unified Presence server.

Step 5   Click **Next**.

Configuring the Local-Side Certificates for the Cisco Presence Federation Proxy

Within an enterprise, setting up a trust relationship is achievable by using self-signed certificates. The supports using self-signed certificates only at this step.

Step 1   In the ASA's Identity Certificate area, click **Generate and Export ASA's Identity Certificate**.

An information dialog box appears indicating that enrollment succeeded. In the Enrollment Status dialog box, click **OK**. The Export certificate dialog box appears.

Note   • If an identity certificate for the ASA has already been created, the button in this area appears as **Export ASA's Identity Certificate** and the Export certificate dialog box immediately appears.

• When using the wizard to configure the Cisco Presence Federation proxy, the wizard only supports installing self-signed certificates.

Step 2   Export the identity certificate generated by the wizard for the ASA. See Exporting an Identity Certificate, page 14-23.

Step 3   Local Unified Presence Server's Certificate area, click **Install Server's Certificate**. The Install Certificate dialog appears.

Step 4   Locate the file containing the Cisco Unified Presence server certificate or paste the certificate details in the dialog box. See Installing a Certificate, page 14-23.

Step 5   Click **Next**.

Note   See the Cisco Unified Presence server documentation for information on how to export the certificate for this server.

Configuring the Remote-Side Certificates for the Cisco Presence Federation Proxy

Establishing a trust relationship across enterprises or across administrative domains is key for federation. Across enterprises you must use a trusted third-party CA (such as, VeriSign). The security appliance obtains a certificate with the FQDN of the Cisco Unified Presence server (certificate impersonation).

For the TLS handshake, the two entities, namely the local entity and a remote entity, could validate the peer certificate via a certificate chain to trusted third-party certificate authorities. The local entity and the remote entity enroll with the CAs. The ASA as the TLS proxy must be trusted by both the local and remote entities. The security appliance is always associated with one of the enterprises. Within that enterprise, the entity and the security appliance authenticate each other by using a self-signed certificate.

To establish a trusted relationship between the security appliance and the remote entity, the security appliance can enroll with the CA on behalf of the Cisco Unified Presence server for the local entity. In the enrollment request, the local entity identity (domain name) is used.

To establish the trust relationship, the security appliance enrolls with the third party CA by using the Cisco Unified Presence server FQDN as if the security appliance is the Cisco Unified Presence server.

Note   If the ASA already has a signed identity certificate, you can skip Step 1 in this procedure and proceed directly to Step 2.

Step 1   In the ASA's Identity Certificate area, click **Generate CSR**. The CSR parameters dialog box appears.

For information about specifying additional parameters for the certificate signing request (CSR), see Generating a Certificate Signing Request (CSR) for a Unified Communications Proxy, page 14-24.

Information dialog boxes appear indicating that the wizard is delivering the settings to the ASA and retrieving the certificate key pair information. The Identity Certificate Request dialog box appears.

For information about saving the CSR that was generated and submitting it to a CA, see Saving the Identity Certificate Request, page 14-25.

Step 2   Click **Install ASA's Identity Certificate**. See Installing the ASA Identity Certificate on the Presence Federation and Cisco Intercompany Media Engine Servers, page 14-26.

Step 3   Click **Remote Server's CA's Certificate**. The Install Certificate dialog box appears. Install the certificate. See Installing a Certificate, page 14-23.

Note   You must install a root CA certificate for each remote entity that communicates with the ASA because different organizations might be using different CAs.

Step 4   Click **Next**.

The wizard completes by displaying a summary of the configuration created for the Presence Federation proxy.

Configuring the UC-IME by using the Unified Communication Wizard

Note   The Unified Communication Wizard is supported for the ASA version 8.3(1) and later.

To configure the Cisco Intercompany Media Engine Proxy by using ASDM, choose **Wizards > Unified Communication Wizard** from the menu. The Unified Communication Wizard opens. From the first page, select the Cisco Intercompany Media Engine Proxy option under the Business-to-Business section and click **Next**.

> **Note**   The Cisco Intercompany Media Engine Proxy does not appear as an option in the Unified Communication Wizard unless the license required for this proxy is installed on the ASA.

When using the wizard to create the Cisco Intercompany Media Engine Proxy, ASDM automatically creates the necessary TLS proxies, enables SIP inspection for Cisco Intercompany Media Engine traffic, generates address translation (static PAT) statements for local Cisco Unified Communications Manager servers, and creates ACLs to allow traffic between the local Cisco Unified Communications Manager servers and the remote servers.

The following steps provide the high-level overview for configuring the Cisco Intercompany Media Engine Proxy:

Step 1   Select the topology of the Cisco Intercompany Media Engine Proxy, namely whether the security appliance is an edge firewall with all Internet traffic flowing through it or whether the security appliance is off the path of the main Internet traffic (referred to as an off-path deployment). See Configuring the Topology for the Cisco Intercompany Media Engine Proxy, page 14-17.

Step 2   Specify private network settings such as the Cisco UCM IP addresses and the ticket settings. See Configuring the Private Network Settings for the Cisco Intercompany Media Engine Proxy, page 14-18.

Step 3   Specify the public network settings. See Configuring the Public Network Settings for the Cisco Intercompany Media Engine Proxy, page 14-20.

Step 4   Specify the media termination address settings of the Cisco UMC. See Configuring the Media Termination Address for Unified Communication Proxies, page 14-10.

Step 5   Configure the local-side certificate management, namely the certificates that are exchanged between the local Cisco Unified Communications Manager servers and the security appliance. See Configuring the Local-Side Certificates for the Cisco Intercompany Media Engine Proxy, page 14-21.

Step 6   Configure the remote-side certificate management, namely the certificates that are exchanged between the remote server and the ASA. This certificate is presented to remote servers so that they can authenticate the ASA as a trusted server. See Configuring the Remote-Side Certificates for the Cisco Intercompany Media Engine Proxy, page 14-22.

The wizard completes by displaying a summary of the configuration created for the Cisco Intercompany Media Engine.

Configuring the Topology for the Cisco Intercompany Media Engine Proxy

Step 1   Select the topology of your ICME deployment by clicking one of the following options:

- All Internet traffic flows through the ASA radio button. This option is also referred to as a basic deployment.
- This ASA is off the path of the regular Internet traffic. This option is also referred to as an off-path deployment.

Step 2   Click **Next**.

Basic Deployment

In a basic deployment, the Cisco Intercompany Media Engine Proxy sits in-line with the Internet firewall such that all Internet traffic traverses the ASA. In this deployment, a single Cisco UCM or a Cisco UCM cluster is centrally deployed within the enterprise, along with a Cisco Intercompany Media Engine server (and perhaps a backup). A single Internet connection traverses the ASA, which is enabled with the Cisco Intercompany Media Engine Proxy.

The ASA sits on the edge of the enterprise and inspects SIP signaling by creating dynamic SIP trunks between enterprises.

Off-path Deployment

In an off path deployment, inbound and outbound Cisco Intercompany Media Engine calls pass through an ASA enabled with the Cisco Intercompany Media Engine Proxy. The ASA is located in the DMZ and configured to support primarily Cisco Intercompany Media Engine. Normal Internet facing traffic does not flow through this ASA.

For all inbound calls, the signaling is directed to the ASA because destined Cisco UCMs are configured with the global IP address on the ASA. For outbound calls, the called party could be any IP address on the Internet; therefore, the ASA is configured with a mapping service that dynamically provides an internal IP address on the ASA for each global IP address of the called party on the Internet.

Cisco UCM sends all outbound calls directly to the mapped internal IP address on the ASA instead of the global IP address of the called party on the Internet. The ASA then forwards the calls to the global IP address of the called party.

> **Note**   When you configure the Cisco Intercompany Media Engine for an off-path deployment, you must ensure that the public IP addresses and ports of the Cisco Unified Communications Manager servers and the public IP address for the media termination address are accessible from the Internet. The summary page of the Unified Communication Wizard reminds you of the requirements.

Configuring the Private Network Settings for the Cisco Intercompany Media Engine Proxy

When configuring the Cisco Intercompany Media Engine Proxy, you specify settings to define the private network topology, such the private network interface, the IP addresses of the Cisco Unified Communications servers, and ticket verification. Additionally, when the Cisco Unified Communications servers are operating in secure mode, you specify the X.509 subject name for the Cisco Intercompany Media Engine Proxy,

The values that you specify in this page generate the following configuration settings for the Cisco Intercompany Media Engine Proxy:

- The list of Cisco Unified Communications servers
- The ticket epoch and password used by the Cisco Intercompany Media Engine Proxy
- For an off-path deployment only, the mapping service on the same interface as the Cisco Unified Communications server

Step 1   To configure the Cisco Intercompany Media Engine Proxy as part of a basic deployment, select the interface that connects to the local Cisco Unified Communications servers.

Or

To configure the Cisco Intercompany Media Engine Proxy as part of an off-path deployment, complete the following steps:

a.   From the Listening Interface drop-down list, choose the interface on which the ASA listens for the mapping requests.

b.   In the Port field, enter a number between 1024 and 65535 as the TCP port on which the ASA listens for the mapping requests. The port number must be 1024 or higher to avoid conflicts with other services on the device, such as Telnet or SSH. By default, the port number is TCP 8060.

c.   From the UC-IME Interface drop-down list, choose the interface that the ASA uses to connect to the remote ASA that is enabled with the Cisco Intercompany Media Engine Proxy.

Note   In a basic and an off-path deployment, all Cisco Unified Communications servers must be on the same interface.

Step 2   In the Unified CM Servers area, the wizard displays the private IP address, public IP address, and security mode of any Cisco Unified Communications server configured on the ASA. If necessary, click **Add** to add a Cisco Unified Communications server. You must include an entry for each Cisco UCM in the cluster with Cisco Intercompany Media Engine that has a SIP trunk enabled.

Step 3   In the Ticket Epoch field, enter a integer from 1-255.

The epoch indicates the number of times that password has changed. When the proxy is configured the first time and a password entered for the first time, enter 1 for the epoch integer. Each time you change the password, increment the epoch to indicate the new password. You must increment the epoch value each time your change the password. Typically, you increment the epoch sequentially; however, the security appliance allows you to choose any value when you update the epoch.

If you change the epoch value, the current password is invalidated and you must enter a new password.

Step 4   In the Ticket Password field, enter a minimum of 10 and a maximum of 64 printable character from the US-ASCII character set. The allowed characters include 0x21 to 0x73 inclusive, and exclude the space character. The ticket password is stored onto flash.

Note   We recommend a password of at least 20 characters. Only one password can be configured at a time.

The epoch and password that you configure on the ASA must match the epoch and password configured on the Cisco Intercompany Media Engine server. See the Cisco Intercompany Media Engine server documentation for information.

Step 5   In the Confirm Password field, reenter the password.

Step 6   In the X.509 Subject Name field, enter the distinguished name (DN) of the local enterprise. The name that you enter must match the name configured for the Cisco Unified Communications servers in the cluster. See the Cisco Unified Communications server documentation for information.

Step 7   Click **Next**.

Cisco ASA Series Firewall ASDM Configuration Guide

Adding a Cisco Unified Communications Manager Server for the UC-IME Proxy

You must include an entry for each Cisco UCM in the cluster with Cisco Intercompany Media Engine Proxy that has a SIP trunk enabled.

Step 1 Enter the private IP address and port number (in the range 5000-6000) for the Cisco UCM server.

Step 2 In the Address Translation area, enter the public IP address for the Cisco UCM server.

Step 3 If necessary, enter the port number for the public IP address by clicking the Translate address and port radio button and entering a number (in the range 5000-6000) in the Port field.

Step 4 In the Security Mode area, click the Secure or Non-secure radio button. Specifying secure for Cisco UCM or Cisco UCM cluster indicates that Cisco UCM or Cisco UCM cluster is initiating TLS.

If you specify that some of the Cisco UCM servers are operating in secure mode, the Unified Communications Wizard includes a step in the proxy configuration to generate certificates for the local-side communication between the ASA and that Cisco UCM server. See Configuring the Local-Side Certificates for the Cisco Intercompany Media Engine Proxy, page 14-21.

Step 5 Click **OK**.

Configuring the Public Network Settings for the Cisco Intercompany Media Engine Proxy

The public network configuration depends on the deployment scenario you selected in the topology step of this wizard. Specifically, when you are configuring the UC-IME proxy as part of an off-path deployment, this step of the wizard displays fields for address translation, requiring that you specify the private IP address for the UC-IME proxy. Specifying this private IP address, translates IP addresses for inbound traffic.

In an off-path deployment, any existing ASA that you have deployed in your environment are not capable of transmitting Cisco Intercompany Media Engine traffic. Therefore, off-path signaling requires that outside addresses translate to an inside (private) IP address. The inside interface address can be used for this mapping service configuration. For the Cisco Intercompany Media Engine Proxy, the ASA creates dynamic mappings for external addresses to the internal IP address.

The values that you specify in this page generate the following configuration settings for the Cisco Intercompany Media Engine Proxy:

- Static PAT for the Cisco Unified Communications servers
- ACLs for traffic between the local and the remote servers

Step 1 In the Configure public network area, choose an interface from the Interface drop-down list.

Step 2 When configuring an off-path deployment, in the Address Translation area, specify whether to use the private IP address for the public network.

Or

Click the Specify IP address radio button and enter an IP address in the field.

Step 3 Click **Next**.

Configuring the Local-Side Certificates for the Cisco Intercompany Media Engine Proxy

Completing this step of the wizard generates a self-signed certificate for the ASA. The server proxy certificate is automatically generated using the subject name provided in an earlier step of this wizard.

The wizard supports using self-signed certificates only.

A trusted relationship between the ASA and the Cisco UMA server can be established with self-signed certificates. The certificates are used by the security appliance and the Cisco UCMs to authenticate each other, respectively, during TLS handshakes.

The ASA's identity certificate is exported, and then needs to be installed on each Cisco Unified Communications Manager (UCM) server in the cluster with the proxy and each identity certificate from the Cisco UCMs need to be installed on the security appliance.

This step in the Unified Communications Wizard only appears when the UC-IME proxy that you are creating has at least one secure Cisco Unified Communications Manager server defined. See Configuring the Topology for the Cisco Intercompany Media Engine Proxy, page 14-17 for information.

Step 1   In the ASA's Identity Certificate area, click **Generate and Export ASA's Identity Certificate**.

An information dialog boxes appear indicating that the enrollment seceded. In the Enrollment Status dialog box, click **OK**. The Export certificate dialog box appears.

Note
- If an identity certificate for the ASA has already been created, the button in this area appears as **Export ASA's Identity Certificate** and the Export certificate dialog box immediately appears.
- When using the wizard to configure the Cisco Intercompany Media Engine Proxy, the wizard only supports installing self-signed certificates.

Step 2   Export the identity certificate generated by the wizard for the ASA. See Exporting an Identity Certificate, page 14-23.

Step 3   In the Local Unified CM's Certificate area, click **Install Local Unified CM's Certificate**. The Install Certificate dialog appears.

Step 4   Locate the file containing the certificate from the Cisco Unified Communications Manager server or paste the certificate details in the dialog box. See Installing a Certificate, page 14-23. You must install the certificate from each Cisco Unified Communications Manager server in the cluster.

Step 5   Click **Next**.

Note   See the Cisco Intercompany Media Engine server documentation for information on how to export the certificate for this server.

Configuring the Remote-Side Certificates for the Cisco Intercompany Media Engine Proxy

Establishing a trust relationship cross enterprises or across administrative domains is key. Cross enterprises you must use a trusted third-party CA (such as, VeriSign). The ASA obtains a certificate with the FQDN of the Cisco Unified Communications Manager server (certificate impersonation).

For the TLS handshake, the two entities could validate the peer certificate via a certificate chain to trusted third-party certificate authorities. Both entities enroll with the CAs. The ASA as the TLS proxy must be trusted by both entities. The ASA is always associated with one of the enterprises. Within that enterprise, the entity and the ASA could authenticate each other via a local CA, or by using self-signed certificates.

To establish a trusted relationship between the ASA and the remote entity, the ASA can enroll with the CA on behalf of the local enterprise. In the enrollment request, the local Cisco UCM identity (domain name) is used.

To establish the trust relationship, the ASA enrolls with the third party CA by using the Cisco Unified Communications Manager server FQDN as if the security appliance is the Cisco UCM.

Note   If the ASA already has a signed identity certificate, you can skip Step 1 in this procedure and proceed directly to Step 3.

Step 1   In the ASA's Identity Certificate area, click **Generate CSR**. The CSR parameters dialog box appears.

For information about specifying additional parameters for the certificate signing request (CSR), see Generating a Certificate Signing Request (CSR) for a Unified Communications Proxy, page 14-24.

Information dialog boxes appear indicating that the wizard is delivering the settings to the ASA and retrieving the certificate key pair information. The Identity Certificate Request dialog box appears.

For information about saving the CSR that was generated and submitting it to a CA, see Saving the Identity Certificate Request, page 14-25.

Step 2   In the ASA's Identity Certificate area, click **Install ASA's Identity Certificate**. Installing the ASA Identity Certificate on the Presence Federation and Cisco Intercompany Media Engine Servers, page 14-26.

Step 3   In the Remote Server's CA's Certificate area, click **Install Remote Server's CA's Certificate**. Installing the root certificates of the CA for the remote servers is necessary so that the ASA can determine that the remote servers are trusted.

The Install Certificate dialog box appears. Install the certificate. See Installing a Certificate, page 14-23.

Note   You must install the root certificates only when the root certificates for the remote servers are received from a CA other than the one that provided the identity certificate for the ASA

Step 4   Click **Next**.

The wizard completes by displaying a summary of the configuration created for the Cisco Intercompany Media Engine.

Working with Certificates in the Unified Communication Wizard

This section includes the following topics:

Exporting an Identity Certificate

The Cisco Mobility Advantage Proxy, Cisco Presence Federation Proxy, or Cisco Intercompany Media Engine Proxy require that you export the ASA identity certificate to install on the Cisco Mobility Advantage server, Cisco Presence Federation server, and Cisco Unified Communications server, respectfully.

You use the wizard to export a self-signed identity certificate. The identity certificate has all associated keys and is in PKCS12 format, which is the public key cryptography standard. When configuring a Unified Communications proxy by using the wizard, you click the Generate and Export ASA's Identify Certificate button while in the local-side or server-side certificate management step of the wizard. The Export certificate dialog box appears.

From the Export certificate dialog box, perform these steps:

Step 1   Enter the name of the PKCS12 format file to use in exporting the certificate configuration. Alternatively, click Browse to display the Export ID Certificate File dialog box to find the file to which you want to export the certificate configuration.

Step 2   Click Export Certificate to export the certificate configuration.

An information dialog box appears informing you that the certificate configuration file has been successfully exported to the location that you specified.

You complete the configuration of the Cisco Mobility Advantage Proxy, Cisco Presence Federation Proxy, or Cisco Intercompany Media Engine Proxy, you must import the generated ASA identify certificate in to the Cisco Mobility Advantage server, Cisco Presence Federation server, and Cisco Unified Communications server, respectfully, depending on which proxy you are configuring.

See the documentation for the for each of these products for information about importing an identity certificate into each.

Installing a Certificate

When configuring certificates for the Phone Proxy, Cisco Mobility Advantage Proxy, the Cisco Presence Federation Proxy, and Cisco Intercompany Media Engine Proxy, you must install the certificates from the Cisco Unified Communications Manager servers, the Cisco Mobility Advantage server, the Cisco

Presence Federation server, and the Cisco Unified Communications Manager servers, respectively, on the ASA. See the documentation for each of these products for information about obtaining the identity certificates from each.

When configuring the Cisco Phone Proxy, if LSC provisioning is required or you have LSC enabled IP phones, you must install the CAPF certificate from the Cisco UCM on the ASA. If the Cisco UCM has more than one CAPF certificate, you must import all of them to the ASA. See Enabling Certificate Authority Proxy Function (CAPF) for IP Phones, page 14-8.

Additionally, when configuring the Cisco Mobility Advantage Proxy, you use the Install Certificate dialog box to install the root certificate received from the certificate authority. The root certificate from the certificate authority is used to sign other certificates. The root certificate is used by the ASA to authenticate your signed identity certificate received from the certificate authority.

Note   When using the wizard to configure the Unified Communications proxies, the wizard only supports installing self-signed certificates.

From the Install Certificate dialog box, perform these steps:

Step 1   Perform one of the following actions:

- To add a certificate configuration from an existing file, click the **Install from a file** radio button (this is the default setting). Enter the path and file name, or click **Browse** to search for the file. Then click **Install Certificate**.

- To enroll manually, click the **Paste certificate in PEM format** radio button. Copy and paste the PEM format (base64 or hexadecimal) certificate into the area provided.

Step 2   Click **Install Certificate**.

An information dialog box appears informing you that the certificate was installed on the ASA successfully.

Generating a Certificate Signing Request (CSR) for a Unified Communications Proxy

When configuring certificates for the Cisco Mobility Advantage Proxy, Cisco Presence Federation Proxy, or Cisco Intercompany Media Engine Proxy, you must generate and identity certificate request for the ASA.

Note   If the ASA already has a signed identity certificate, you do not need to generate a CSR and can proceed directly to installing this certificate on the ASA. See Installing the ASA Identity Certificate on the Mobility Advantage Server, page 14-26 and Installing the ASA Identity Certificate on the Presence Federation and Cisco Intercompany Media Engine Servers, page 14-26 for the steps to install the identity certificate.

The identify certificate that you receive is presented to the following entities for each of the Unified Communication Proxies:

- Unified Mobile Communicator clients for the Cisco Mobility Advantage Proxy

- Remote Presence Federation servers for the Cisco Presence Federation Proxy
- The remote ASAfor the Cisco Intercompany Media Engine Proxy

Before generating the CSR, you can enter additional parameters.

When configuring a Unified Communications proxy by using the wizard, you click the Generate CSR button while in the client-side or remote-side certificate management step of the wizard. The CSR Parameters dialog box appears.

In the CSR Parameters dialog box, perform the following steps:

Step 1   From the Key Pair Size drop-down list, choose the size required for you certificate.

The key size that you select depends on the level of security that you want to configure and on any limitations imposed by the CA from which you are obtaining the certificate. The larger the number that you select, the higher the security level will be for the certificate. Most CAs recommend 2048 for the key modulus size; however, GoDaddy requires a key modulus size of 2048.

Step 2   (Cisco Intercompany Media Engine Proxy only) In the CN field, enter the domain name used by your enterprise or network. The subject DN you configure for the Cisco Intercompany Media Engine Proxy must match the domain name that set in the local Cisco Unified Communications Manager server.

> **Note**   For the Cisco Mobility Advantage Proxy and Cisco Presence Federation Proxy, the wizard provides the common name (CN), which is the FQDN of the Cisco Mobility Advantage server or Cisco Unified Presence server, respectively.

Step 3   In the Additional DN Attributes field, enter an attribute.

Or

Click **Select** to display the Additional DN Attributes dialog box.

a.   In the Additional DN Attributes dialog box, choose an attribute from the drop-down list.

b.   Enter a value for the attribute.

c.   Click Add. The attribute appears in the list.

d.   Click OK to return to the CSR Parameters dialog box.

The value you added appears in the Additional DN Attributes field in the CSR Parameters dialog box.

Step 4   Click **OK**.

Saving the Identity Certificate Request

After successfully generating the identity certificate request for one of the Unified Communications proxies, the Identity Certificate Request dialog box appears and prompts you to save the request.

Step 1   In the Save CSR to File field, enter the CSR file name and path; for example, c:\asa-csr.txt.

Step 2   Click **OK**. An information dialog box appears indicating the CSR was saved successfully.

Step 3   Click **OK** to close the dialog and return to the wizard.

Submit the CSR to the certificate authority (CA), for example, by pasting the CSR text into the CSR enrollment page on the CA website.

When the CA returns the signed identity certificate, rerun the Unified Communications Wizard. From the client-side or remote-side certificate management step of the wizard, click **Install ASA's Identity Certificate**. See Installing the ASA Identity Certificate on the Mobility Advantage Server, page 14-26 and Installing the ASA Identity Certificate on the Presence Federation and Cisco Intercompany Media Engine Servers, page 14-26 for the steps to install the identity certificate.

Installing the ASA Identity Certificate on the Mobility Advantage Server

When configuring certificates for the Cisco Mobility Advantage Proxy, you must install the ASA identity certificate on the Cisco Mobility Advantage server.

Typically, a certificate authority returns two certificates: your signed identity certificate and the certificate authority's certificate (referred to as the root certificate). However, some certificate authorities (for example, VeriSign) might also send you an intermediate certificate.

The root certificate from the certificate authority is used to sign other certificates. The root certificate is used by the ASA to authenticate your signed identity certificate received from the certificate authority.

If the certificate authority provided an intermediate certificate, you must enter the certificate text in the Intermediate Certificate (If Applicable) area of the Install ASA's Identity Certificate dialog box.

For the Cisco Mobility Advantage Proxy, you install the root certificate in another dialog box. See Installing a Certificate, page 14-23 for the steps to install the root certificate.

Step 1   In the Intermediate Certificate (If Applicable) area, perform on of the following actions:

- To add a certificate configuration from an existing file, click the **Install from a file** radio button (this is the default setting). Enter the path and file name, or click **Browse** to search for the file. Then click **Install Certificate**.

- To enroll manually, click the **Paste the certificate data in base-64 format** radio button. Copy and paste the PEM format (base64 or hexadecimal) certificate into the area provided.

Step 2   In the ASA's Identity Certificate area, perform on of the following actions:

- To add a certificate configuration from an existing file, click the **Install from a file** radio button (this is the default setting). Enter the path and file name, or click **Browse** to search for the file. Then click **Install Certificate**.

- To enroll manually, click the **Paste the certificate data in base-64 format** radio button. Copy and paste the PEM format (base64 or hexadecimal) certificate into the area provided.

Step 3   Click **Install Certificate**.

Installing the ASA Identity Certificate on the Presence Federation and Cisco Intercompany Media Engine Servers

When configuring certificates for the Cisco Presence Federation Proxy and Cisco Intercompany Media Engine Proxy, you must install the ASA identity certificate and the root certificate on the Cisco Presence Federation server and Cisco Intercompany Media Engine server, respectively.

Typically, a certificate authority returns two certificates: your signed identity certificate and the certificate authority's certificate (referred to as the root certificate). The root certificate from the certificate authority is used to sign other certificates. The root certificate is used by the ASA to authenticate your signed identity certificate received from the certificate authority.

Step 1   In the Root CA's Certificate area, perform on of the following actions:

- To add a certificate configuration from an existing file, click the **Install from a file** radio button (this is the default setting). Enter the path and file name, or click **Browse** to search for the file. Then click **Install Certificate**.

- To enroll manually, click the **Paste the certificate data in base-64 format** radio button. Copy and paste the PEM format (base64 or hexadecimal) certificate into the area provided.

Step 2   In the ASA's Identity Certificate area, perform on of the following actions:

- To add a certificate configuration from an existing file, click the **Install from a file** radio button (this is the default setting). Enter the path and file name, or click **Browse** to search for the file. Then click **Install Certificate**.

- To enroll manually, click the **Paste the certificate data in base-64 format** radio button. Copy and paste the PEM format (base64 or hexadecimal) certificate into the area provided.

Step 3   Click **Install Certificate**.

Cisco Phone Proxy

This chapter describes how to configure the ASA for Cisco Phone Proxy feature.

This chapter includes the following sections:

Information About the Cisco Phone Proxy

The Cisco Phone Proxy on the ASA bridges IP telephony between the corporate IP telephony network and the Internet in a secure manner by forcing data from remote phones on an untrusted network to be encrypted.

Phone Proxy Functionality

Telecommuters can connect their IP phones to the corporate IP telephony network over the Internet securely via the phone proxy without the need to connect over a VPN tunnel as illustrated by Figure 15-1.

Figure 15-1   Phone Proxy Secure Deployment

The phone proxy supports a Cisco UCM cluster in mixed mode or nonsecure mode. Regardless of the cluster mode, the remote phones that are capable of encryption are always forced to be in encrypted mode. TLS (signaling) and SRTP (media) are always terminated on the ASA. The ASA can also perform NAT, open pinholes for the media, and apply inspection policies for the SCCP and SIP protocols. In a nonsecure cluster mode or a mixed mode where the phones are configured as nonsecure, the phone proxy behaves in the following ways:

- The TLS connections from the phones are terminated on the ASA and a TCP connection is initiated to the Cisco UCM.
- SRTP sent from external IP phones to the internal network IP phone via the ASA is converted to RTP.

In a mixed mode cluster where the internal IP phones are configured as authenticated, the TLS connection is not converted to TCP to the Cisco UCM but the SRTP is converted to RTP.

In a mixed mode cluster where the internal IP phone is configured as encrypted, the TLS connection remains a TLS connection to the Cisco UCM and the SRTP from the remote phone remains SRTP to the internal IP phone.

Since the main purpose of the phone proxy is to make the phone behave securely while making calls to a nonsecure cluster, the phone proxy performs the following major functions:

- Creates the certificate trust list (CTL) file, which is used to perform certificate based authentication with remote phones.
- Modifies the IP phone configuration file when it is requested via TFTP, changes security fields from nonsecure to secure, and signs all files sent to the phone. These modifications secure remote phones by forcing the phones to perform encrypted signaling and media.
- Terminates TLS signaling from the phone and initiates TCP or TLS to Cisco UCM
- Inserts itself into the media path by modifying the Skinny and SIP signaling messages.
- Terminates SRTP and initiates RTP/SRTP to the called party.

Note   As an alternative to authenticating remote IP phones through the TLS handshake, you can configure authentication via LSC provisioning. With LSC provisioning you create a password for each remote IP phone user and each user enters the password on the remote IP phones to retrieve the LSC.

Because using LSC provisioning to authenticate remote IP phones requires the IP phones first register in nonsecure mode, Cisco recommends LSC provisioning be done inside the corporate network before giving the IP phones to end-users. Otherwise, having the IP phones register in nonsecure mode requires the Administrator to open the nonsecure signaling port for SIP and SCCP on the ASA.

See also the Cisco Unified Communications Manager Security Guide for information on Using the Certificate Authority Proxy Function (CAPF) to install a locally significant certificate (LSC).

Supported Cisco UCM and IP Phones for the Phone Proxy

Cisco Unified Communications Manager

The following release of the Cisco Unified Communications Manager are supported with the phone proxy:

- Cisco Unified CallManager Version 4.*x*
- Cisco Unified CallManager Version 5.0
- Cisco Unified CallManager Version 5.1
- Cisco Unified Communications Manager 6.1
- Cisco Unified Communications Manager 7.0
- Cisco Unified Communications Manager 8.0

Cisco Unified IP Phones

The phone proxy supports these IP phone features:

- Enterprise features like conference calls on remote phones connected through the phone proxy
- XML services

The following IP phones in the Cisco Unified IP Phones 7900 Series are supported with the phone proxy:

- Cisco Unified IP Phone 7975
- Cisco Unified IP Phone 7971
- Cisco Unified IP Phone 7970
- Cisco Unified IP Phone 7965
- Cisco Unified IP Phone 7962
- Cisco Unified IP Phone 7961
- Cisco Unified IP Phone 7961G-GE
- Cisco Unified IP Phone 7960 (SCCP protocol support only)
- Cisco Unified IP Phone 7945
- Cisco Unified IP Phone 7942
- Cisco Unified IP Phone 7941

- Cisco Unified IP Phone 7941G-GE
- Cisco Unified IP Phone 7940 (SCCP protocol support only)
- Cisco Unified Wireless IP Phone 7921
- Cisco Unified Wireless IP Phone 7925

> **Note**   To support Cisco Unified Wireless IP Phone 7925, you must also configure MIC or LSC on the IP phone so that it properly works with the phone proxy.

- CIPC for softphones (CIPC versions with Authenticated mode only)

> **Note**   The Cisco IP Communicator is supported with the phone proxy VLAN Traversal in authenticated TLS mode. We do not recommend it for remote access because SRTP/TLS is not supported currently on the Cisco IP Communicator.

> **Note**   The ASA supports inspection of traffic from Cisco IP Phones running SCCP protocol version 19 and earlier.

Licensing Requirements for the Phone Proxy

The Cisco Phone Proxy feature supported by the ASA require a Unified Communications Proxy license.

The following table shows the Unified Communications Proxy license details by platform:

> **Note**   This feature is not available on No Payload Encryption models.

Model	License Requirement[1]
ASA 5505	Base License and Security Plus License: 2 sessions. *Optional license: 24 sessions.*
ASA 5512-X	Base License or Security Plus License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5515-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5525-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, or 1000 sessions.*
ASA 5545-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, or 2000 sessions.*
ASA 5555-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*

Model	License Requirement[1]
ASA 5585-X with SSP-10	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5585-X with SSP-20, -40, or -60	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*
ASASM	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*
ASAv with 1 Virtual CPU	Standard and Premium Licenses: 250 sessions.
ASAv with 4 Virtual CPUs	Standard and Premium Licenses: 1000 sessions.

1. The following applications use TLS proxy sessions for their connections. Each TLS proxy session used by these applications (and only these applications) is counted against the UC license limit:
   - Phone Proxy
   - Presence Federation Proxy
   - Encrypted Voice Inspection

   Other applications that use TLS proxy sessions do not count towards the UC limit, for example, Mobility Advantage Proxy (which does not require a license) and IME (which requires a separate IME license).

   Some UC applications might use multiple sessions for a connection. For example, if you configure a phone with a primary and backup Cisco Unified Communications Manager, there are 2 TLS proxy connections, so 2 UC Proxy sessions are used.

   You independently set the TLS proxy limit using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. When you apply a UC license that is higher than the default TLS proxy limit, the ASA automatically sets the TLS proxy limit to match the UC limit. The TLS proxy limit takes precedence over the UC license limit; if you set the TLS proxy limit to be less than the UC license, then you cannot use all of the sessions in your UC license.

   Note: For license part numbers ending in "K8" (for example, licenses under 250 users), TLS proxy sessions are limited to 1000. For license part numbers ending in "K9" (for example, licenses 250 users or larger), the TLS proxy limit depends on the configuration, up to the model limit. K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted.

   Note: If you clear the configuration, then the TLS proxy limit is set to the default for your model; if this default is lower than the UC license limit, then you see an error message to use the to raise the limit again (in ASDM, use the **TLS Proxy** pane). If you use failover and use **File > Save Running Configuration to Standby Unit** on the primary unit to force a configuration synchronization, the **clear configure all** command is generated on the secondary unit automatically, so you may see the warning message on the secondary unit. Because the configuration synchronization restores the TLS proxy limit set on the primary unit, you can ignore the warning.

   You might also use SRTP encryption sessions for your connections:
   - For K8 licenses, SRTP sessions are limited to 250.
   - For K9 licenses, there is not limit.

   Note: Only calls that require encryption/decryption for media are counted towards the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count towards the limit.

For more information about licensing, see the general operations configuration guide.

Prerequisites for the Phone Proxy

This section contains the following topics:

Media Termination Instance Prerequisites

The ASA must have a media termination instance that meets the following criteria:

- You must configure one media termination for each phone proxy on the ASA. Multiple media termination instances on the ASA are not supported.

- For the media termination instance, you can configure a global media-termination address for all interfaces or configure a media-termination address for different interfaces. However, you cannot use a global media-termination address and media-termination addresses configured for each interface at the same time.

- If you configure a media termination address for multiple interfaces, you must configure an address on each interface that the ASA uses when communicating with IP phones.

  For example, if you had three interfaces on the ASA (one internal interface and two external interfaces) and only one of the external interfaces were used to communicate with IP phones, you would configure two media termination addresses: one on the internal interface and one on the external interface that communicated with the IP phones.

- Only one media-termination address can be configured per interface.

- The IP addresses are publicly routable addresses that are unused IP addresses within the address range on that interface.

- The IP address on an interface cannot be the same address as that interface on the ASA.

- The IP addresses cannot overlap with existing static NAT pools or NAT rules.

- The IP addresses cannot be the same as the Cisco UCM or TFTP server IP address.

- For IP phones behind a router or gateway, you must also meet this prerequisite. On the router or gateway, add routes to the media termination address on the ASA interface that the IP phones communicate with so that the phone can reach the media termination address.

Certificates from the Cisco UCM

Import the following certificates which are stored on the Cisco UCM. These certificates are required by the ASA for the phone proxy.

- Cisco_Manufacturing_CA
- CAP-RTP-001
- CAP-RTP-002
- CAPF certificate (Optional)

If LSC provisioning is required or you have LSC enabled IP phones, you must import the CAPF certificate from the Cisco UCM. If the Cisco UCM has more than one CAPF certificate, you must import all of them to the ASA.

Note   You can configure LSC provisioning for additional end-user authentication. See the Cisco Unified Communications Manager configuration guide for information.

For example, the CA Manufacturer certificate is required by the phone proxy to validate the IP phone certificate.

DNS Lookup Prerequisites

- If you have an fully qualified domain name (FQDN) configured for the Cisco UCM rather than an IP address, you must configure and enable DNS lookup on the ASA.
- After configuring the DNS lookup, make sure that the ASA can ping the Cisco UCM with the configured FQDN.
- You must configure DNS lookup when you have a CAPF service enabled and the Cisco UCM is not running on the Publisher but the Publisher is configured with a FQDN instead of an IP address.

Cisco Unified Communications Manager Prerequisites

- The TFTP server must reside on the same interface as the Cisco UCM.
- The Cisco UCM can be on a private network on the inside but you need to have a static mapping for the Cisco UCM on the ASA to a public routable address.
- If NAT is required for Cisco UCM, it must be configured on the ASA, not on the existing firewall.

ACL Rules

If the phone proxy is deployed behind an existing firewall, access-list rules to permit signaling, TFTP requests, and media traffic to the phone proxy must be configured.

If NAT is configured for the TFTP server or Cisco UCMs, the translated "global" address must be used in the ACLs.

Table 15-1 lists the ports that are required to be configured on the existing firewall:

Table 15-1   Port Configuration Requirements

Address	Port	Protocol	Description
Media Termination	1024-65535	UDP	Allow incoming SRTP
TFTP Server	69	UDP	Allow incoming TFTP
Cisco UCM	2443	TCP	Allow incoming secure SCCP

Table 15-1   Port Configuration Requirements

Address	Port	Protocol	Description
Cisco UCM	5061	TCP	Allow incoming secure SIP
CAPF Service (on Cisco UCM)	3804	TCP	Allow CAPF service for LSC provisioning

✎

Note   All these ports are configurable on the Cisco UCM, except for TFTP. These are the default values and should be modified if they are modified on the Cisco UCM. For example, 3804 is the default port for the CAPF Service. This default value should be modified if it is modified on the Cisco UCM.

NAT and PAT Prerequisites

NAT Prerequisites

- If NAT is configured for the TFTP server, the NAT configuration must be configured prior to configuring the TFTP Server for the phone proxy.

- If NAT is configured for the TFTP server or Cisco UCMs, the translated "global" address must be used in the ACLs.

PAT Prerequisites

- When the Skinny inspection global port is configured to use a non-default port, then you must configure the nonsecure port as the global_sccp_port+443.

  Therefore, if *global_sccp_port* is 7000, then the global secure SCCP port is 7443. Reconfiguring the port might be necessary when the phone proxy deployment has more than one Cisco UCM and they must share the interface IP address or a global IP address.

  ✎

  Note   Both PAT configurations—for the nonsecure and secure ports—must be configured.

- When the IP phones must contact the CAPF on the Cisco UCM and the Cisco UCM is configured with static PAT (LCS provisioning is required), you must configure static PAT for the default CAPF port 3804.

Prerequisites for IP Phones on Multiple Interfaces

When IP phones reside on multiple interfaces, the phone proxy configuration must have the correct IP address set for the Cisco UCM in the CTL file.

See the following example topology for information about how to correctly set the IP address:

```
phones --- (dmz)-----|
                     |----- ASA PP --- (outside Internet) --- phones
phones --- (inside)--|
```

In this example topology, the following IP address are set:

- Cisco UCM on the inside interface is set to 10.0.0.5
- The DMZ network is 192.168.1.0/24
- The inside network is 10.0.0.0/24

The Cisco UCM is mapped with different global IP addresses from DMZ > outside and inside interfaces > outside interface.

In the CTL file, the Cisco UCM must have two entries because of the two different IP addresses. For example, if the static statements for the Cisco UCM are as follows:

```
object network obj-10.0.0.5-01
    host 10.0.0.5
    nat (inside,outside) static 209.165.202.129
object network obj-10.0.0.5-02
    host 10.0.0.5
    nat (inside,dmz) static 198.168.1.2
```

There must be two CTL file record entries for the Cisco UCM:

```
record-entry cucm trustpoint cucm_in_to_out address 209.165.202.129
record-entry cucm trustpoint cucm_in_to_dmz address 192.168.1.2
```

7960 and 7940 IP Phones Support

- An LSC must be installed on these IP phones because they do not come pre installed with a MIC. Install the LSC on each phone before using them with the phone proxy to avoid opening the nonsecure SCCP port for the IP phones to register in nonsecure mode with the Cisco UCM.

  See the following document for the steps to install an LSC on IP phones:

  http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/security/7_0_1/secugd/secucapf.html#wp1093518

  Note If an IP phone already has an LSC installed on it from a different Cisco UCM cluster, delete the LSC from the different cluster and install an LSC from the current Cisco UCM cluster.

  Note You can configure LSC provisioning for additional end-user authentication. See the Cisco Unified Communications Manager configuration guide for information.

- The CAPF certificate must be imported onto the ASA.
- The CTL file created on the ASA must be created with a CAPF record-entry.
- The phone must be configured to use only the SCCP protocol because the SIP protocol does not support encryption on these IP phones.
- If LSC provisioning is done via the phone proxy, you must add an ACL to allow the IP phones to register with the Cisco UCM on the nonsecure port 2000.

Cisco IP Communicator Prerequisites

To configure Cisco IP Communicator (CIPC) with the phone proxy, you must meet the following prerequisites:

- Go to Configuration > Firewall > Unified Communications > Phone Proxy and select the "Enable CICP security mode authentication" check box under the Call Manager and Phone Settings area.

- Create an ACL to allow CIPC to register with the Cisco UCM in nonsecure mode.

- Configure null-sha1 as one of the SSL encryption ciphers.

Current versions of Cisco IP Communicator (CIPC) support authenticated mode and perform TLS signaling but not voice encryption.

Because CIPC requires an LSC to perform the TLS handshake, CIPC needs to register with the Cisco UCM in nonsecure mode using cleartext signaling. To allow the CIPC to register, create an ACL that allows the CIPC to connect to the Cisco UCM on the nonsecure SIP/SCCP signalling ports (5060/2000).

Note   You can configure LSC provisioning for additional end-user authentication. See the Cisco Unified Communications Manager configuration guide for information.

CIPC uses a different cipher when doing the TLS handshake and requires the null-sha1 cipher and SSL encryption be configured. To add the null-shal cipher, use the show run all ssl command to see the output for the ssl encryption command and add null-shal to the end of the SSL encryption list.

Note   When used with CIPC, the phone proxy does not support end-users resetting their device name in CIPC (Preferences > Network tab > Use this Device Name field) or Administrators resetting the device name in Cisco Unified CM Administration console (Device menu > Phone Configuration > Device Name field). To function with the phone proxy, the CIPC configuration file must be in the format: SEP<mac_address>.cnf.xml. If the device name does not follow this format (SEP<mac_address>), CIPC cannot retrieve its configuration file from Cisco UMC via the phone proxy and CIPC will not function.

Prerequisites for Rate Limiting TFTP Requests

In a remote access scenario, we recommend that you configure rate limiting of TFTP requests because any IP phone connecting through the Internet is allowed to send TFTP requests to the TFTP server.

To configure rate limiting of TFTP requests, configure the **police** command in the Modular Policy Framework. See the command reference for information about using the **police** command.

Policing is a way of ensuring that no traffic exceeds the maximum rate (in bits/second) that you configure, thus ensuring that no one traffic flow can take over the entire resource. When traffic exceeds the maximum rate, the ASA drops the excess traffic. Policing also sets the largest single burst of traffic allowed.

Rate Limiting Configuration Example

The following example describes how you configure rate limiting for TFTP requests by using the **police** command and the Modular Policy Framework.

Begin by determining the conformance rate that is required for the phone proxy. To determine the conformance rate, use the following formula:

```
X * Y * 8
```

Where

X = requests per second

Y = size of each packet, which includes the L2, L3, and L4 plus the payload

Therefore, if a rate of 300 TFTP requests/second is required, then the conformance rate would be calculated as follows:

```
300 requests/second * 80 bytes * 8 = 192000
```

To control which hosts can ping the media termination address, create an ICMP rule. Go to Configuration > Device Management > Management Access > ICMP and click the Add button.

End-User Phone Provisioning

The phone proxy is a transparent proxy with respect to the TFTP and signaling transactions. If NAT is not configured for the Cisco UCM TFTP server, then the IP phones need to be configured with the Cisco UCM cluster TFTP server address.

If NAT is configured for the Cisco UCM TFTP server, then the Cisco UCM TFTP server global address is configured as the TFTP server address on the IP phones.

Ways to Deploy IP Phones to End Users

In both options, deploying a remote IP phone behind a commercial Cable/DSL router with NAT capabilities is supported.

Option 1 (Recommended)

Stage the IP phones at corporate headquarters before sending them to the end users:

- The phones register inside the network. IT ensures there are no issues with the phone configurations, image downloads, and registration.
- If Cisco UCM cluster was in mixed mode, the CTL file should be erased before sending the phone to the end user.

Advantages of this option are:

- Easier to troubleshoot and isolate problems with the network or phone proxy because you know whether the phone is registered and working with the Cisco UCM.
- Better user experience because the phone does not have to download firmware from over a broadband connection, which can be slow and require the user to wait for a longer time.

Option 2

Send the IP phone to the end user. When using option 2, the user must be provided instructions to change the settings on phones with the appropriate Cisco UCM and TFTP server IP address.

Note As an alternative to authenticating remote IP phones through the TLS handshake, you can configure authentication via LSC provisioning. With LSC provisioning you create a password for each remote IP phone user and each user enters the password on the remote IP phones to retrieve the LSC.

Because using LSC provisioning to authenticate remote IP phones requires the IP phones first register in nonsecure mode, Cisco recommends LSC provisioning be done inside the corporate network before giving the IP phones to end-users. Otherwise, having the IP phones register in nonsecure mode requires the Administrator to open the nonsecure signaling port for SIP and SCCP on the ASA.

See also the Cisco Unified Communications Manager Security Guide for information on Using the Certificate Authority Proxy Function (CAPF) to install a locally significant certificate (LSC).

Phone Proxy Guidelines and Limitations

This section includes the following topics:

General Guidelines and Limitations

The phone proxy has the following general limitations:

- Only one phone proxy instance can be configured on the ASA by using the **phone-proxy** command. See the command reference for information about the **phone-proxy** command. See also Creating the Phone Proxy Instance, page 15-17.

- The phone proxy only supports one Cisco UCM cluster. See Creating the CTL File, page 15-14 for the steps to configure the Cisco UCM cluster for the phone proxy.

- The phone proxy is not supported when the ASA is running in transparent mode or multiple context mode.

- When a remote IP phone calls an invalid internal or external extension, the phone proxy does not support playing the annunciator message from the Cisco UCM. Instead, the remote IP phone plays a fast busy signal instead of the annunciator message "Your call cannot be completed ..." However, when an internal IP phone dials in invalid extension, the annunciator messages plays "Your call cannot be completed ..."

- Packets from phones connecting to the phone proxy over a VPN tunnel are not inspected by the ASA inspection engines.

- The phone proxy does not support IP phones sending Real-Time Control Protocol (RTCP) packets through the ASA. Disable RTCP packets in the Cisco Unified CM Administration console from the Phone Configuration page. See your Cisco Unified Communications Manager (CallManager) documentation for information about setting this configuration option.

- When used with CIPC, the phone proxy does not support end-users resetting their device name in CIPC (Preferences > Network tab > Use this Device Name field) or Administrators resetting the device name in Cisco Unified CM Administration console (Device menu > Phone Configuration > Device Name field). To function with the phone proxy, the CIPC configuration file must be in the format: SEP<mac_address>.cnf.xml. If the device name does not follow this format (SEP<mac_address>), CIPC cannot retrieve its configuration file from Cisco UMC via the phone proxy and CIPC will not function.

- The phone proxy does not support IP phones sending SCCP video messages using Cisco VT Advantage because SCCP video messages do not support SRTP keys.

- For mixed-mode clusters, the phone proxy does not support the Cisco Unified Call Manager using TFTP to send encrypted configuration files to IP phones through the ASA.

- Multiple IP phones behind one NAT device must be configured to use the same security mode.

When the phone proxy is configured for a mixed-mode cluster and multiple IP phones are behind one NAT device and registering through the phone proxy, all the SIP and SCCP IP phones must be configured as authenticated or encrypted, or all as non-secure on the Unified Call Manager.

For example, if there are four IP phones behind one NAT device where two IP phones are configured using SIP and two IP phones are configured using SCCP, the following configurations on the Unified Call Manager are acceptable:

- Two SIP IP phones: one IP phone in authenticated mode and one in encrypted mode, both in authenticated mode, or both in encrypted mode

  Two SCCP IP phones: one IP phone in authenticated mode and one in encrypted mode, both in authenticated mode, or both in encrypted mode

- Two SIP IP phones: both in non-secure mode

  Two SCCP IP phones: one IP phone in authenticated mode and one in encrypted mode, both in authenticated mode, both in encrypted mode

- Two SIP IP phones: one IP phone in authenticated mode and one in encrypted mode, both in authenticated mode, both in encrypted mode

  Two SCCP IP phones: both in non-secure mode

This limitation results from the way the application-redirect rules (rules that convert TLS to TCP) are created for the IP phones.

Media Termination Address Guidelines and Limitations

The phone proxy has the following limitations relating to configuring the media-termination address:

- When configuring the media-termination address, the phone proxy does not support having internal IP phones (IP phones on the inside network) being on a different network interface from the Cisco UCM unless the IP phones are forced to use the non-secure Security mode.

  When internal IP phones are on a different network interface than the Cisco UCM, the IP phones signalling sessions still go through ASA; however, the IP phone traffic does not go through the phone proxy. Therefore, Cisco recommends that you deploy internal IP phones on the same network interface as the Cisco UMC.

  If the Cisco UMC and the internal IP phones must be on different network interfaces, you must add routes for the internal IP phones to access the network interface of the media-termination address where Cisco UMC resides.

  When the phone proxy is configured to use a global media-termination address, all IP phones see the same global address, which is a public routable address.

- If you decide to configure a media-termination address on interfaces (rather than using a global interface), you must configure a media-termination address on at least two interfaces (the inside and an outside interface) before applying the phone-proxy service policy. Otherwise, you will receive an error message when enabling the Phone Proxy with SIP and Skinny Inspection.

- The phone proxy can use only one type of media termination instance at a time; for example, you can configure a global media-termination address for all interfaces or configure a media-termination address for different interfaces. However, you cannot use a global media-termination address and media-termination addresses configured for each interface at the same time.

Configuring the Phone Proxy

This section includes the following topics:

Task Flow for Configuring the Phone Proxy

Note   This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Configuring the Phone Proxy requires the following steps:

Step 1: Create the CTL file. See Creating the CTL File, page 15-14.

Step 2: Create the TLS Proxy instance to handle the encrypted signaling. See Adding a TLS Proxy Instance, page 16-8.

Step 3: Create the Phone Proxy instance. See Creating the Phone Proxy Instance, page 15-17.

Step 4: Configure the media termination address for the Phone Proxy. See Creating the Media Termination Instance, page 15-16.

Note   Before you enable SIP and Skinny inspection for the Phone Proxy (which is done by applying the Phone Proxy to a service policy rule), the Phone Proxy must have an MTA instance, TLS Proxy, and CTL file assigned to it before the Phone Proxy can be applied to a service policy. Additionally, once a Phone Proxy is applied to a service policy rule, the Phone Proxy cannot be changed or removed.

Step 5: Enable the Phone Proxy with SIP and Skinny inspection. See SIP Inspection, page 10-20 and Skinny (SCCP) Inspection, page 10-32.

Creating the CTL File

Create a Certificate Trust List (CTL) file that is required by the Phone Proxy. Specify the certificates needed by creating a new CTL file or by specifying the path of an exiting CTL file to parse from Flash memory.

Create trustpoints and generate certificates for each entity in the network (CUCM, CUCM and TFTP, TFTP server, CAPF) that the IP phones must trust. The certificates are used in creating the CTL file. You need to create trustpoints for each CUCM (primary and secondary if a secondary CUCM is used) and TFTP server in the network. The trustpoints need to be in the CTL file for the phones to trust the CUCM.

Create the CTL File that will be presented to the IP phones during the TFTP. The address must be the translated or global address of the TFTP server or CUCM if NAT is configured.

When the file is created, it creates an internal trustpoint used by the Phone Proxy to sign the TFTP files. The trustpoint is named **_internal_PP**_*ctl-instance_filename*.

> **Note** When a CTL file instance is assigned to the Phone Proxy, you cannot modify it in the CTL File pane and the pane is disabled. To modify a CTL File that is assigned to the Phone Proxy, go to the Phone Proxy pane (Configuration > Firewall > Unified Communications > Phone Proxy), and deselect the Use the Certificate Trust List File generated by the CTL instance check box.

Use the Create a Certificate Trust List (CTL) File pane to create a CTL file for the Phone Proxy. This pane creates the CTL file that is presented to the IP phones during the TFTP handshake with the ASA. For a detailed overview of the CTL file used by the Phone Proxy, see Creating the CTL File, page 15-14.

The Create a Certificate Trust List (CTL) File pane is used to configure the attributes for generating the CTL file. The name of the CTL file instance is generated by the ASDM. When the user tries to edit the CTL file instance configuration, the ASDM automatically generates the **shutdown** CLI command first and the **no shutdown** CLI command as the last command.

This pane is available from the Configuration > Firewall > Unified Communications > CTL File pane.

Step 1 Open the Configuration > Firewall > Unified Communications > CTL File pane.

Step 2 Check the Enable Certificate Trust List File check box to enable the feature.

Step 3 To specify the CTL file to use for the Phone Proxy, perform one of the following:

- If there is an existing CTL file available, download the CTL file to Flash memory by using the File Management Tool in the ASDM Tools menu. Select the Use certificates present in the CTL stored in flash radio button and specify the CTL file name and path in the text box.

  Use an existing CTL file to install the trustpoints for each entity in the network (CUCM, CUCM and TFTP, TFTP server, CAPF) that the IP phones must trust. If you have an existing CTL file that contains the correct IP addresses of the entities (namely, the IP address that the IP phones use for the CUCM or TFTP servers), you can be use it to create a new CTL file. Store a copy of the existing CTL file to Flash memory and rename it something other than CTLFile.tlv

- If there is no existing CTL file available, select Create new CTL file radio button.

  Add Record entries for each entity in the network such as CUCM, TFTP, and CUCM-TFTP option by clicking **Add**. The Add Record Entry dialog box opens. See Adding or Editing a Record Entry in a CTL File, page 15-15.

Step 4 Specify the number SAST certificate tokens required. The default is 2. maximum allowed is 5.

Because the Phone Proxy generates the CTL file, it needs to create the System Administrator Security Token (SAST) key to sign the CTL file itself. This key can be generated on the ASA. A SAST is created as a self-signed certificate. Typically, a CTL file contains more than one SAST. In case a SAST is not recoverable, the other one can be used to sign the file later.

Step 5 Click **Apply** to save the CTL file configuration settings.

Adding or Editing a Record Entry in a CTL File

> **Note** This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Use the Add/Edit Record Entry dialog box to specify the trustpoints to be used for the creation of the CTL file.

> **Note** You can edit an entry in the CTL file by using the Edit Record Entry dialog box; however, changing a setting in this dialog box does not change related settings for the phone proxy. For example, editing the IP address for the CUCM or TFTP servers in this dialog changes the setting only in the CTL file and does not change the actual addresses of those servers or update the address translations required by the phone proxy.
>
> To modify CTL file settings, we strongly recommend you re-run the Unified Communications Wizard to edit CTL file settings and ensure proper synchronization with all phone proxy settings.

Add additional record-entry configurations for each entity that is required in the CTL file.

Step 1 Open the Configuration > Firewall > Unified Communications > CTL File pane.

Step 2 Check the Enable Certificate Trust List File check box to enable the feature.

Step 3 In the Type field, specify the type of trustpoint to create:

- cucm: Specifies the role of this trustpoint to be CCM. Multiple CCM trustpoints can be configured.
- cucm-tftp: Specifies the role of this trustpoint to be CCM+TFTP. Multiple CCM+TFTP trustpoints can be configured.
- tftp: Specifies the role of this trustpoint to be TFTP. Multiple TFTP trustpoints can be configured.
- capf: Specifies the role of this trustpoint to be CAPF. Only one CAPF trustpoint can be configured.

Step 4 In the Host field, specify the IP address of the trustpoint. The IP address you specify must be the global address of the TFTP server or CUCM if NAT is configured. The global IP address is the IP address as seen by the IP phones because it will be the IP address used for the CTL record for the trustpoint.

Step 5 In the Certificate field, specify the Identity Certificate for the record entry in the CTL file. You can create a new Identity Certificate by clicking **Manage**. The Manage Identify Certificates dialog box opens. See the general operations configuration guide.

You can add an Identity Certificate by generating a self-signed certificate, obtaining the certificate through SCEP enrollment, or by importing a certificate in PKCS-12 format. Choose the best option based on the requirements for configuring the CTL file.

Step 6 (Optional) In the Domain Name field, specify the domain name of the trustpoint used to create the DNS field for the trustpoint. This is appended to the Common Name field of the Subject DN to create the DNS Name. The domain name should be configured when the FQDN is not configured for the trustpoint. Only one domain-name can be specified.

> **Note** If you are using domain names for your CUCM and TFTP server, you must configure DNS lookup on the ASA. Add an entry for each of the outside interfaces on the ASA into your DNS server, if such entries are not already present. Each ASA outside IP address should have a DNS entry associated with it for lookups. These DNS entries must also be enabled for Reverse Lookup. Additionally, define your DNS server IP address on the ASA; for example: `dns name-server 10.2.3.4` (IP address of your DNS server).

Creating the Media Termination Instance

Create the media termination instance that you will use in the phone proxy.

The media termination address you configure must meet the requirements as described in Media Termination Instance Prerequisites, page 15-6.

> **Note** In versions before 8.2(1), you configured one media-termination address (MTA) on the outside interface of the adaptive security appliance where the remote Cisco IP phones were located. In Version 8.2(1) and later, you can configure a global media-termination address for all interfaces or configure a media-termination address for different interfaces.
>
> As a result of this enhancement, the old configuration has been deprecated. You can continue to use the old configuration if desired. However, if you need to change the configuration at all, only the new configuration method is accepted; you cannot later restore the old configuration. If you need to maintain downgrade compatibility, you should keep the old configuration as is.

Step 1 Open the Configuration > Firewall > Unified Communications > Media Termination Address pane.

Step 2 Check the Enable Media Termination Address check box to enable the feature.

Step 3 In the Media Termination Address Settings area, specify whether to configure a media-termination address (MTA) per interface or to configure a global MTA. You can configure a global media-termination address for all interfaces or configure a media-termination address for different interfaces.

- To configure an MTA per interface, click the Configure MTA per Interface radio button and click the **Add** button. In the dialog box that appears, specify the interface name and enter an IP address or hostname.

  If you configure a media termination address for multiple interfaces, you must configure an address on each interface that the ASA uses when communicating with IP phones. The IP addresses are publicly routable addresses that are unused IP addresses within the address range on that interface.

  See Media Termination Instance Prerequisites, page 15-6 for the complete list of requirements that you must follow when creating the media termination instance and configuring the media termination addresses.

- To configure a global MTA, click the Configure global MTA on interface radio button and enter the IP address in the text box. See Media Termination Instance Prerequisites, page 15-6 for the complete list of requirements that you must follow when configuring a global media termination address.

Step 4 Specify the minimum and maximum values for the RTP port range for the media termination instance. The minimum port and the maximum port can be a value from 1024 to 65535.

Step 5 Click **Apply** to save the media termination address configuration settings.

Creating the Phone Proxy Instance

Create the phone proxy instance. To have a fully functional phone proxy, you must also complete additional tasks, such as creating the MTA and enabling SIP and SCCP (Skinny) inspection. See Task Flow for Configuring the Phone Proxy, page 15-14 for the complete list of tasks.

Prerequisites

You must have already created the CTL file and TLS proxy instance for the phone proxy.

See Creating the CTL File, page 15-14 and Adding a TLS Proxy Instance, page 16-8.

> **Note** This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Use the Configure Phone Proxy pane to add a Phone Proxy.

This pane is available from the Configuration > Firewall > Unified Communications > Phone Proxy pane.

Step 1 Open the Configuration > Firewall > Unified Communications > Phone Proxy pane.

Step 2 Check the Enable Phone Proxy check box to enable the feature.

Step 3 Check the Apply MTA instance to Phone Proxy check box to add the media termination address to the Phone Proxy instance. You must have a media termination address instance configured. The configured address is added to the Phone Proxy instance.

> **Note** To configure the media termination address, click the Configure MTA button. The Media Termination Address dialog box appears. Once you click the Add MTA instance to Phone Proxy check box, the media termination address instance cannot be modified and the button changes to View MTA Configuration. To change the media termination address, uncheck the Add MTA instance to Phone Proxy check box.

Step 4 If necessary, add a TFTP server for the Phone Proxy. To add a new TFTP server for the Phone Proxy, click **Add**. The Add TFTP Server dialog box opens. See Adding or Editing the TFTP Server for a Phone Proxy, page 15-20.

> **Note** The TFTP server must reside on the same interface as the Cisco Unified Call Manager. Additionally, If NAT is configured for the TFTP server, the NAT configuration must be configured prior to configuring the specifying the TFTP server while creating the Phone Proxy instance.

Step 5 Specify the CTL File to use for the Phone Proxy by doing one of the following:

- To use an existing CTL File, check the Use the Certificate Trust List File generated by the CTL instance check box.

- To create a new CTL file for the Phone Proxy, click the link Generate Certificate Trust List File. The Create a Certificate Trust List (CTL) File pane opens. See Creating the CTL File, page 15-14.

Step 6 To specify the security mode of the CUCM cluster, click one of the following options in the CUCM Cluster Mode field:

- Non-secure—Specifies the cluster mode to be in nonsecure mode when configuring the Phone Proxy feature.

- Mixed—Specifies the cluster mode to be in mixed mode when configuring the Phone Proxy feature.

Step 7 To configure the idle timeout after which the secure-phone entry is removed from the Phone Proxy database (the default is 5 minutes), enter a value in the format *hh*:*mm*:*ss*.

Since secure phones always request a CTL file upon bootup, the Phone Proxy creates a database that marks the phone as secure. The entries in the secure phone database are removed after a specified configured timeout. The entry timestamp is updated for each registration refresh the Phone Proxy receives for SIP phones and KeepAlives for SCCP phones.

Specify a value that is greater than the maximum timeout value for SCCP KeepAlives and SIP Register refresh. For example, if the SCCP KeepAlives are configured for 1 minute intervals and the SIP Register Refresh is configured for 3 minutes, configure this timeout value greater than 3 minutes.

Step 8 To preserve Call Manager configuration on the IP phones, check the Preserve the Call Manager's configuration on the phone... check box. When this check box is uncheck, the following service settings are disabled on the IP phones:

- PC Port
- Gratuitous ARP
- Voice VLAN access
- Web Access
- Span to PC Port

Step 9 To force Cisco IP Communicator (CIPC) softphones to operate in authenticated mode when CIPC softphones are deployed in a voice and data VLAN scenario, check the Enable CIPC security mode authentication check box.

Because CIPC requires an LSC to perform the TLS handshake, CIPC needs to register with the CUCM in nonsecure mode using cleartext signaling. To allow the CIPC to register, create an ACL that allows the CIPC to connect to the CUCM on the nonsecure SIP/SCCP signalling ports (5060/2000).

CIPC uses a different cipher when doing the TLS handshake and requires the null-sha1 cipher and SSL encryption be configured. To add the null-sha1 cipher, go to Configuration > Device Management > Advanced > SSL Settings > Encryption section. Select the null-sha1 SSL encryption type and add it to the Available Algorithms.

Current versions of Cisco IP Communicator (CIPC) support authenticated mode and perform TLS signaling but not voice encryption.

Step 10 To configure an HTTP proxy for the Phone Proxy feature that is written into the IP phone's configuration file under the <proxyServerURL> tag, do the following:

- **a.** Check the Configure a http-proxy which would be written into the phone's config file... check box.
- **b.** In the IP Address field, type the IP address of the HTTP proxy and the listening port of the HTTP proxy.

  The IP address you enter should be the global IP address based on where the IP phone and HTTP proxy server is located. You can enter a hostname in the IP Address field when that hostname can be resolved to an IP address by the ASA (for example, DNS lookup is configured) because the ASA will resolve the hostname to an IP address. If a port is not specified, the default will be 8080.

- **c.** In the Interface field, select the interface on which the HTTP proxy resides on the ASA.

Setting the proxy server configuration option for the Phone Proxy allows for an HTTP proxy on the DMZ or external network in which all the IP phone URLs are directed to the proxy server for services on the phones. This setting accommodates nonsecure HTTP traffic, which is not allowed back into the corporate network.

Step 11 Click **Apply** to save the Phone Proxy configuration settings.

Note After creating the Phone Proxy instance, you enable it with SIP and Skinny inspection. See SIP Inspection, page 10-20 and Skinny (SCCP) Inspection, page 10-32.

However, before you enable SIP and Skinny inspection for the Phone Proxy (which is done by applying

the Phone Proxy to a service policy rule), the Phone Proxy must have an MTA instance, TLS Proxy, and CTL file assigned to it before the Phone Proxy can be applied to a service policy. Additionally, once a Phone Proxy is applied to a service policy rule, the Phone Proxy cannot be changed or removed.

Adding or Editing the TFTP Server for a Phone Proxy

Note   This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Note   You can edit the TFTP server setting by using the Edit TFTP Server dialog box; however, changing a setting in this dialog box does not change related settings for the phone proxy. For example, editing the IP address for the TFTP server in this dialog does not change the setting in the CTL file and does not update the address translations required by the phone proxy.
To modify TFTP server settings, we strongly recommend you re-run the Unified Communications Wizard to ensure proper synchronization with all phone proxy settings.

Step 1   Open the Configuration > Firewall > Unified Communications > Phone Proxy pane.

Step 2   Check the Enable Phone Proxy check box to enable the feature.

Step 3   To add or edit the TFTP Server information for the phone proxy, click the **Add** or **Edit** button. The Add/Edit TFTP Server dialog box appears.

Use the Add/Edit TFTP Server dialog box to specify the IP address of the TFTP server and the interface on which the TFTP server resides.

The Phone Proxy must have at least one CUCM TFTP server configured. Up to five TFTP servers can be configured for the Phone Proxy.

The TFTP server is assumed to be behind the firewall on the trusted network; therefore, the Phone Proxy intercepts the requests between the IP phones and TFTP server.

Note   If NAT is configured for the TFTP server, the NAT configuration must be configured prior to specifying the TFTP server while creating the Phone Proxy instance.

Step 4   In the TFTP Server IP Address field, specify the address of the TFTP server. Create the TFTP server using the actual internal IP address.

Step 5   (Optional) In the Port field, specify the port the TFTP server is listening in on for the TFTP requests. This should be configured if it is not the default TFTP port 69.

Step 6   In the Interface field, specify the interface on which the TFTP server resides. The TFTP server must reside on the same interface as the Cisco Unified Call Manager (CUCM).

Step 7   Click OK to apply the settings.

Configuring Linksys Routers with UDP Port Forwarding for the Phone Proxy

When IP phones are behind a NAT-capable router, the router can be configured to forward the UDP ports to the IP address of the IP phone. Specifically, configure the router for UDP port forwarding when an IP phone is failing during TFTP requests and the failure is due to the router dropping incoming TFTP data packets. Configure the router to enable UDP port forwarding on port 69 to the IP phone.

As an alternative of explicit UDP forwarding, some Cable/DSL routers require you to designate the IP phone as a DMZ host. For Cable/DSL routers, this host is a special host that receives all incoming connections from the public network.

When configuring the phone proxy, there is no functional difference between an IP phone that has UDP ports explicitly forwarded or an IP phone designated as a DMZ host. The choice is entirely dependent upon the capabilities and preference of the end user.

Configuring Your Router

Your firewall/router needs to be configured to forward a range of UDP ports to the IP phone. This will allow the IP phone to receive audio when you make/receive calls.

> **Note**   Different Cable/DSL routers have different procedures for this configuration. Furthermore most NAT-capable routers will only allow a given port range to be forwarded to a single IP address

The configuration of each brand/model of firewall/router is different, but the task is the same. For specific instructions for your brand and model of router, please contact the manufacturer's website.

Linksys Routers

Step 1   From your web browser, connect to the router administrative web page. For Linksys, this is typically something like `http://192.168.1.1`.

Step 2   Click Applications & Gaming or the Port Forwarding tab (whichever is present on your router).

Step 3   Locate the table containing the port forwarding data and add an entry containing the following values:

Table 15-2   Port Forwarding Values to Add to Router

Application	Start	End	Protocol	IP Address	Enabled
IP phone	1024	65535	UDP	*Phone IP address*	**Checked**
TFTP	69	69	UDP	*Phone IP address*	**Checked**

Step 4   Click Save Settings. Port forwarding is configured.

Feature History for the Phone Proxy

Table 15-3 lists the release history for this feature.

Table 15-3   Feature History for Cisco Phone Proxy

Feature Name	Releases	Feature Information
Cisco Phone Proxy	8.0(4)	The phone proxy feature was introduced. The Phone Proxy feature was accessible in ASDM by choosing the following options: Configuration > Firewall > Advanced > Encrypted Traffic Inspection > Phone Proxy pane
NAT for the media termination address	8.1(2)	The Media Termination fields were removed from the Phone Proxy pane and added to the Media Termination pane: Configuration > Firewall > Advanced > Encrypted Traffic Inspection > Media Termination Address pane

CHAPTER **16**

TLS Proxy for Encrypted Voice Inspection

This chapter describes how to configure the ASA for the TLS Proxy for Encrypted Voice Inspection feature.

This chapter includes the following sections:

Information about the TLS Proxy for Encrypted Voice Inspection

End-to-end encryption often leaves network security appliances "blind" to media and signaling traffic, which can compromise access control and threat prevention security functions. This lack of visibility can result in a lack of interoperability between the firewall functions and the encrypted voice, leaving businesses unable to satisfy both of their key security requirements.

The ASA is able to intercept and decrypt encrypted signaling from Cisco encrypted endpoints to the Cisco Unified Communications Manager (Cisco UCM), and apply the required threat protection and access control. It can also ensure confidentiality by re-encrypting the traffic onto the Cisco UCM servers.

Typically, the ASA TLS Proxy functionality is deployed in campus unified communications network. This solution is ideal for deployments that utilize end to end encryption and firewalls to protect Unified Communications Manager servers.

The security appliance in Figure 16-1 serves as a proxy for both client and server, with Cisco IP Phone and Cisco UCM interaction.

Figure 16-1   TLS Proxy Flow
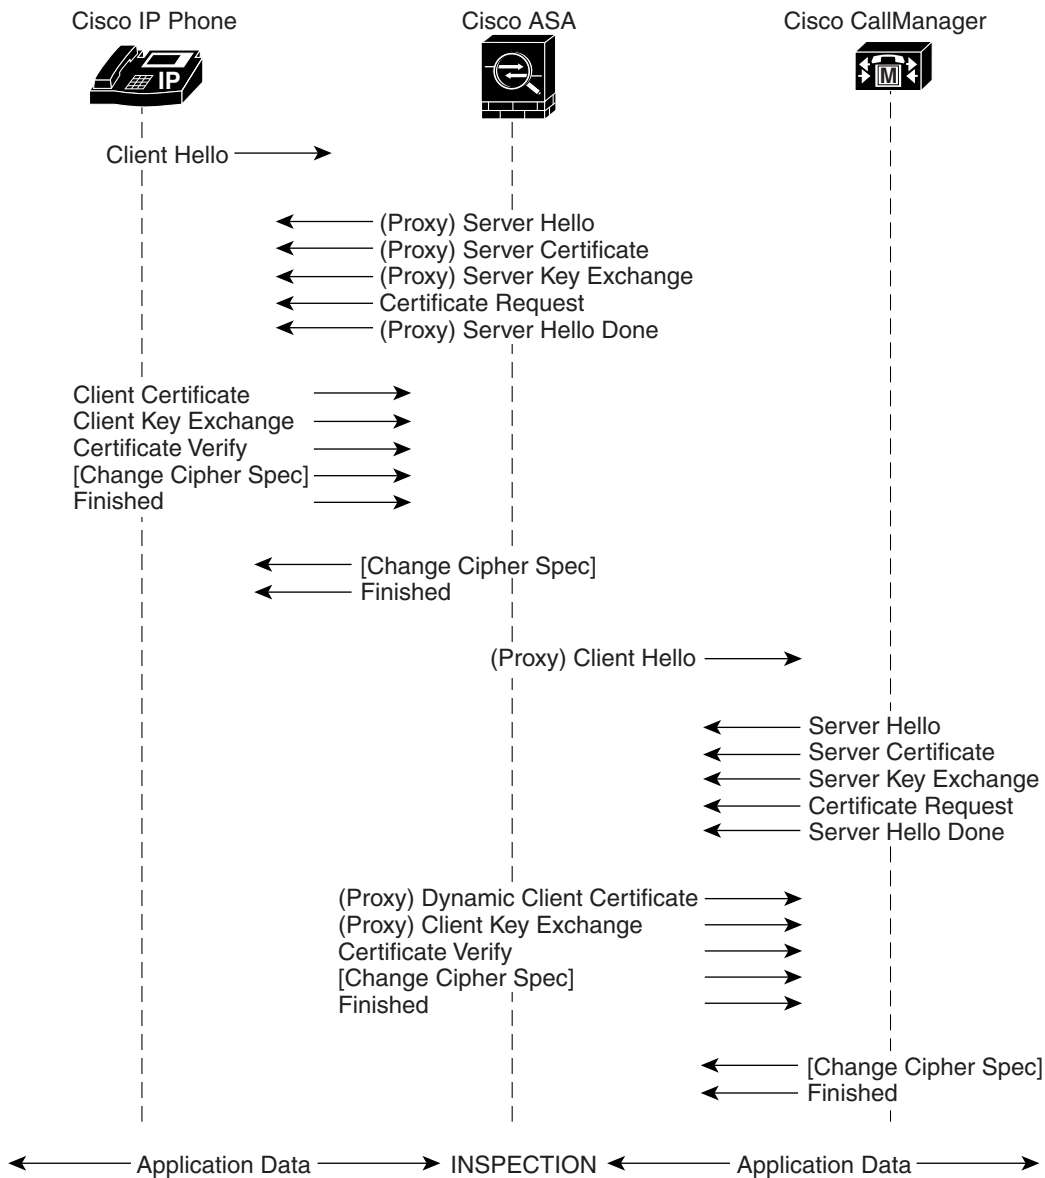
Decryption and Inspection of Unified Communications Encrypted Signaling

With encrypted voice inspection, the security appliance decrypts, inspects and modifies (as needed, for example, performing NAT fixup), and re-encrypts voice signaling traffic while all of the existing VoIP inspection functions for Skinny and SIP protocols are preserved. Once voice signaling is decrypted, the plaintext signaling message is passed to the existing inspection engines.

The security appliance acts as a TLS proxy between the Cisco IP Phone and Cisco UCM. The proxy is transparent for the voice calls between the phone and theCisco UCM. Cisco IP Phones download a Certificate Trust List from the Cisco UCM before registration which contains identities (certificates) of the devices that the phone should trust, such as TFTP servers and Cisco UCM servers. To support server

proxy, the CTL file must contain the certificate that the security appliance creates for the Cisco UCMs. To proxy calls on behalf of the Cisco IP Phone, the security appliance presents a certificate that the Cisco UCM can verify, which is a Local Dynamic Certificate for the phone, issued by the certificate authority on the security appliance.

TLS proxy is supported by the Cisco Unified CallManager Release 5.1 and later. You should be familiar with the security features of the Cisco UCM. For background and detailed description of Cisco UCM security, see the Cisco Unified CallManager document:

http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/5_0/sec_vir/ae/sec504/index.htm

TLS proxy applies to the encryption layer and must be configured with an application layer protocol inspection. You should be familiar with the inspection features on the ASA, especially Skinny and SIP inspection.

Supported Cisco UCM and IP Phones for the TLS Proxy

Cisco Unified Communications Manager

The following releases of the Cisco Unified Communications Manager are supported with the TLS proxy:

- Cisco Unified CallManager Version 4.*x*
- Cisco Unified CallManager Version 5.0
- Cisco Unified CallManager Version 5.1
- Cisco Unified Communications Manager 6.1
- Cisco Unified Communications Manager 7.0
- Cisco Unified Communications Manager 8.0

Cisco Unified IP Phones

The following IP phones in the Cisco Unified IP Phones 7900 Series are supported with the TLS proxy:

- Cisco Unified IP Phone 7985
- Cisco Unified IP Phone 7975
- Cisco Unified IP Phone 7971
- Cisco Unified IP Phone 7970
- Cisco Unified IP Phone 7965
- Cisco Unified IP Phone 7962
- Cisco Unified IP Phone 7961
- Cisco Unified IP Phone 7961G-GE
- Cisco Unified IP Phone 7960
- Cisco Unified IP Phone 7945
- Cisco Unified IP Phone 7942
- Cisco Unified IP Phone 7941
- Cisco Unified IP Phone 7941G-GE
- Cisco Unified IP Phone 7940
- Cisco Unified Wireless IP Phone 7921

- Cisco Unified Wireless IP Phone 7925
- Cisco IP Communicator (CIPC) for softphones

Licensing for the TLS Proxy

The TLS proxy for encrypted voice inspection feature supported by the ASA require a Unified Communications Proxy license.

The following table shows the Unified Communications Proxy license details by platform:

Note   This feature is not available on No Payload Encryption models.

Model	License Requirement[1]
ASA 5505	Base License and Security Plus License: 2 sessions. *Optional license: 24 sessions.*
ASA 5512-X	Base License or Security Plus License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5515-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5525-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, or 1000 sessions.*
ASA 5545-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, or 2000 sessions.*
ASA 5555-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5585-X with SSP-10	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5585-X with SSP-20, -40, or -60	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*
ASASM	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*
ASAv with 1 Virtual CPU	Standard and Premium Licenses: 250 sessions.
ASAv with 4 Virtual CPUs	Standard and Premium Licenses: 1000 sessions.

1. The following applications use TLS proxy sessions for their connections. Each TLS proxy session used by these applications (and only these applications) is counted against the UC license limit:
    - Phone Proxy
    - Presence Federation Proxy
    - Encrypted Voice Inspection

    Other applications that use TLS proxy sessions do not count towards the UC limit, for example, Mobility Advantage Proxy (which does not require a license) and IME (which requires a separate IME license).

    Some UC applications might use multiple sessions for a connection. For example, if you configure a phone with a primary and backup Cisco Unified Communications Manager, there are 2 TLS proxy connections, so 2 UC Proxy sessions are used.

    You independently set the TLS proxy limit using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. When you apply a UC license that is higher than the default TLS proxy limit, the ASA automatically sets the TLS proxy limit to match the UC limit. The TLS proxy limit takes precedence over the UC license limit; if you set the TLS proxy limit to be less than the UC license, then you cannot use all of the sessions in your UC license.

    Note: For license part numbers ending in "K8" (for example, licenses under 250 users), TLS proxy sessions are limited to 1000. For license part numbers ending in "K9" (for example, licenses 250 users or larger), the TLS proxy limit depends on the configuration, up to the model limit. K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted.

    Note: If you clear the configuration, then the TLS proxy limit is set to the default for your model; if this default is lower than the UC license limit, then you see an error message to use the to raise the limit again (in ASDM, use the **TLS Proxy** pane). If you use failover and use **File > Save Running Configuration to Standby Unit** on the primary unit to force a configuration synchronization, the **clear configure all** command is generated on the secondary unit automatically, so you may see the warning message on the secondary unit. Because the configuration synchronization restores the TLS proxy limit set on the primary unit, you can ignore the warning.

    You might also use SRTP encryption sessions for your connections:
    - For K8 licenses, SRTP sessions are limited to 250.
    - For K9 licenses, there is not limit.

    Note: Only calls that require encryption/decryption for media are counted towards the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count towards the limit.

    Table 16-1 shows the default and maximum TLS session details by platform.

Table 16-1   Default and Maximum TLS Sessions on the Security Appliance

Security Appliance Platform	Default TLS Sessions	Maximum TLS Sessions
ASA 5505	10	80

    For more information about licensing, see the general operations configuration guide.

Prerequisites for the TLS Proxy for Encrypted Voice Inspection

Before configuring TLS proxy, the following prerequisites are required:

- You must set clock on the security appliance before configuring TLS proxy. To set the clock manually and display clock, use the **clock set** and **show clock** commands. We recommend that the security appliance use the same NTP server as the Cisco Unified CallManager cluster. TLS handshake may fail due to certificate validation failure if clock is out of sync between the security appliance and the Cisco Unified CallManager server.

- 3DES-AES license is needed to interoperate with the Cisco Unified CallManager. AES is the default cipher used by the Cisco Unified CallManager and Cisco IP Phone.

- Import the following certificates which are stored on the Cisco UCM. These certificates are required by the ASA for the phone proxy.

- Cisco_Manufacturing_CA

- CAP-RTP-001

- CAP-RTP-002

- CAPF certificate (Optional)

  If LSC provisioning is required or you have LSC enabled IP phones, you must import the CAPF certificate from the Cisco UCM. If the Cisco UCM has more than one CAPF certificate, you must import all of them to the ASA.

See Chapter 15, "Cisco Phone Proxy." For example, the CA Manufacturer certificate is required by the phone proxy to validate the IP phone certificate.

Configuring the TLS Proxy for Encrypted Voice Inspection

This section includes the following topics:

CTL Provider

Use the CTL Provider option to configure Certificate Trust List provider service.

The CTL Provider pane lets you define and configure Certificate Trust List provider service to enable inspection of encrypted traffic.

Fields

- CTL Provider Name—Lists the CTL Provider name.

- Client Details—Lists the name and IP address of the client.

  - Interface Name—Lists the defined interface name.

  - IP Address—Lists the defined interface IP address.

- Certificate Name—Lists the certificate to be exported.

- Add—Adds a CTL Provider.

- Edit—Edits a CTL Provider.

- Delete—Deletes a CTL Provider.

Add/Edit CTL Provider

The Add/Edit CTL Provider dialog box lets you define the parameters for the CTL Provider.

Fields

- CTL Provider Name—Specifies the CTL Provider name.
- Certificate to be Exported—Specifies the certificate to be exported to the client.
  - Certificate Name—Specifies the name of the certificate to be exported to the client.
  - Manage—Manages identity certificates.
- Client Details—Specifies the clients allowed to connect.
  - Client to be Added—Specifies the client interface and IP address to add to the client list.

    Interface—Specifies client interface.

    IP Address—Specifies the client IP address.

    Add—Adds the new client to the client list.

    Delete—Deletes the selected client from the client list.
- More Options—Specifies the available and active algorithms to be announced or matched during the TLS handshake.
  - Parse the CTL file provided by the CTL Client and install trustpoints—Trustpoints installed by this option have names prefixed with "_internal_CTL_." If disabled, each Call Manager server and CAPF certificate must be manually imported and installed.
  - Port Number—Specifies the port to which the CTL provider listens. The port must be the same as the one listened to by the CallManager servers in the cluster (as configured under Enterprise Parameters on the CallManager administration page). The default is 2444.
  - Authentication—Specifies the username and password that the client authenticates with the provider.

    Username—Client username.

    Password—Client password.

    Confirm Password—Client password.

Configure TLS Proxy Pane

> **Note** This feature is not supported for the Adaptive Security Appliance version 8.1.2.

You can configure the TLS Proxy from the Configuration > Firewall > Unified Communications > TLS Proxy pane.

Configuring a TLS Proxy lets you use the TLS Proxy to enable inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and enable the ASA for the Cisco Unified Communications features:

- TLS Proxy for the Cisco Unified Presence Server (CUPS), part of Presence Federation
- TLS Proxy for the Cisco Unified Mobility Advantage (CUMA) server, part of Mobile Advantage
- Phone Proxy

Fields

- TLS Proxy Name—Lists the TLS Proxy name.

- Server Proxy Certificate—Lists the trustpoint, which is either self-signed or enrolled with a certificate server.

- Local Dynamic Certificate Issuer—Lists the local certificate authority to issue client or server dynamic certificates.

- Client Proxy Certificate—Lists the proxy certificate for the TLS client. The ASA uses the client proxy certificate to authenticate the TLS client during the handshake between the proxy and the TLS client. The certificate can be either self-signed, enrolled with a certificate authority, or issued by the third party.

- Add—Adds a TLS Proxy by launching the Add TLS Proxy Instance Wizard. See Adding a TLS Proxy Instance, page 16-8 for the steps to create a TLS Proxy instance.

- Edit—Edits a TLS Proxy. The fields in the Edit panel area identical to the fields displayed when you add a TLS Proxy instance. See Edit TLS Proxy Instance – Server Configuration, page 16-12 and Edit TLS Proxy Instance – Client Configuration, page 16-13.

- Delete—Deletes a TLS Proxy.

- Maximum Sessions—Lets you specify the maximum number of TLS Proxy sessions to support.

    - Specify the maximum number of TLS Proxy sessions that the ASA needs to support.

    - Maximum number of sessions—The minimum is 1. The maximum is dependent on the platform.

> **Note**   The maximum number of sessions is global to all TLS proxy sessions.

Adding a TLS Proxy Instance

> **Note**   This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Use the Add TLS Proxy Instance Wizard to add a TLS Proxy to enable inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and to support the Cisco Unified Communications features on the ASA.

This wizard is available from the Configuration > Firewall > Unified Communications > TLS Proxy pane.

Step 1   Open the Configuration > Firewall > Unified Communications > TLS Proxy pane.

Step 2   To add a new TLS Proxy Instance, click **Add**.

The Add TLS Proxy Instance Wizard opens.

Step 3   In the TLS Proxy Name field, type the TLS Proxy name.

Step 4   Click **Next**.

The Add TLS Proxy Instance Wizard – Server Configuration dialog box opens. In this step of the wizard, configure the server proxy parameters for original TLS Server—the Cisco Unified Call Manager (CUCM) server, the Cisco Unified Presence Server (CUPS), or the Cisco Unified Mobility Advantage (CUMA) server. See Add TLS Proxy Instance Wizard – Server Configuration, page 16-9.

After configuring the server proxy parameters, the wizard guides you through configuring client proxy parameters (see Add TLS Proxy Instance Wizard – Client Configuration, page 16-10) and provides instructions on the steps to complete outside the ASDM to make the TLS Proxy fully functional (see Add TLS Proxy Instance Wizard – Other Steps, page 16-11).

Add TLS Proxy Instance Wizard – Server Configuration

Note   This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Use the Add TLS Proxy Instance Wizard to add a TLS Proxy to enable inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and to support the Cisco Unified Communications features on the ASA.

The Add TLS Proxy Instance Wizard is available from the Configuration > Firewall > Unified Communications > TLS Proxy pane.

Step 1   Complete the first step of the Add TLS Proxy Instance Wizard. See Adding a TLS Proxy Instance, page 16-8.

The Add TLS Proxy Instance Wizard – Server Configuration dialog box opens.

Step 2   Specify the server proxy certificate by doing one of the following:

- To add a new certificate, click **Manage**. The Manage Identify Certificates dialog box opens.

    When the Phone Proxy is operating in a mixed-mode CUCM cluster, you must import the CUCM certificate by clicking **Add** in the Manage Identify Certificates dialog box.

- To select an existing certificate, select one from the drop-down list.

    When you are configuring the TLS Proxy for the Phone Proxy, select the certificate that has a filename beginning with **_internal_PP_**. When you create the CTL file for the Phone Proxy, the ASA, creates an internal trustpoint used by the Phone Proxy to sign the TFTP files. The trustpoint is named **_internal_PP_**_ctl-instance_filename_.

The server proxy certificate is used to specify the trustpoint to present during the TLS handshake. The trustpoint can be self-signed or enrolled locally with the certificate service on the proxy. For example, for the Phone Proxy, the server proxy certificate is used by the Phone Proxy during the handshake with the IP phones.

Step 3   To install the TLS server certificate in the ASA trust store, so that the ASA can authenticate the TLS server during TLS handshake between the proxy and the TLS server, click **Install TLS Server's Certificate**.

The Manage CA Certificates dialog box opens. Click **Add** to open the Install Certificate dialog box.

When you are configuring the TLS Proxy for the Phone Proxy, click **Install TLS Server's Certificate** and install the Cisco Unified Call Manager (CUCM) certificate so that the proxy can authenticate the IP phones on behalf of the CUCM server.

Step 4   To require the ASA to present a certificate and authenticate the TLS client during TLS handshake, check the Enable client authentication during TLS Proxy handshake check box.

When adding a TLS Proxy Instance for Mobile Advantage (the CUMC client and CUMA server), disable the check box when the client is incapable of sending a client certificate.

Step 5   Click **Next**.

The Add TLS Proxy Instance Wizard – Client Configuration dialog box opens. In this step of the wizard, configure the client proxy parameters for original TLS Client—the CUMC client for Mobile Advantage, CUP or MS LCS/OCS client for Presence Federation, or the IP phone for the Phone Proxy. See Add TLS Proxy Instance Wizard – Client Configuration, page 16-10.

After configuring the client proxy parameters, the wizard provides instructions on the steps to complete outside the ASDM to make the TLS Proxy fully functional (see Add TLS Proxy Instance Wizard – Other Steps, page 16-11).

Add TLS Proxy Instance Wizard – Client Configuration

Note   This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Use the Add TLS Proxy Instance Wizard to add a TLS Proxy to enable inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and to support the Cisco Unified Communications features on the ASA.

This wizard is available from the Configuration > Firewall > Unified Communications > TLS Proxy pane.

Step 1   Complete the first two steps of the Add TLS Proxy Instance Wizard. See Adding a TLS Proxy Instance, page 16-8 and Add TLS Proxy Instance Wizard – Client Configuration, page 16-10.

The Add TLS Proxy Instance Wizard – Client Configuration dialog box opens.

Step 2   To specify a client proxy certificate to use for the TLS Proxy, perform the following. Select this option when the client proxy certificate is being used between two servers; for example, when configuring the TLS Proxy for Presence Federation, which uses the Cisco Unified Presence Server (CUPS), both the TLS client and TLS server are both servers.

   a.   Check the Specify the proxy certificate for the TLS Client... check box.

   b.   Select a certificate from the drop-down list.

      Or

      To create a new client proxy certificate, click **Manage**. The Manage Identify Certificates dialog box opens.

Note   When you are configuring the TLS Proxy for the Phone Proxy and it is using the mixed security mode for the CUCM cluster, you must configure the LDC Issuer. The LDC Issuer lists the local certificate authority to issue client or server dynamic certificates.

Step 3   To specify an LDC Issuer to use for the TLS Proxy, perform the following. When you select and configure the LDC Issuer option, the ASA acts as the certificate authority and issues certificates to TLS clients.

   a.   Click the Specify the internal Certificate Authority to sign the local dynamic certificate for phones... check box.

b. Click the Certificates radio button and select a self-signed certificate from the drop-down list or click **Manage** to create a new LDC Issuer. The Manage Identify Certificates dialog box opens.

Or

Click the Certificate Authority radio button to specify a Certificate Authority (CA) server. When you specify a CA server, it needs to be created and enabled in the ASA. To create and enable the CA server, click **Manage**. The Edit CA Server Settings dialog box opens.

> ✎
> **Note**   To make configuration changes after the local certificate authority has been configured for the first time, disable the local certificate authority.

c. In the Key-Pair Name field, select a key pair from the drop-list. The list contains the already defined RSA key pair used by client dynamic certificates. To see the key pair details, including generation time, usage, modulus size, and key data, click **Show**.

Or

To create a new key pair, click **New**. The Add Key Pair dialog box opens.

Step 4   In the Security Algorithms area, specify the available and active algorithms to be announced or matched during the TLS handshake.

- Available Algorithms—Lists the available algorithms to be announced or matched during the TLS handshake: des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, and null-sha1.

  Add—Adds the selected algorithm to the active list.

  Remove—Removes the selected algorithm from the active list.

- Active Algorithms—Lists the active algorithms to be announced or matched during the TLS handshake: des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, and null-sha1. For client proxy (acting as a TLS client to the server), the user-defined algorithms replace the original ones from the hello message for asymmetric encryption method between the two TLS legs. For example, the leg between the proxy and Call Manager may be NULL cipher to offload the Call Manager.

  Move Up—Moves an algorithm up in the list.

  Move Down—Moves an algorithm down in the list.

Step 5   Click **Next**.

The Add TLS Proxy Instance Wizard – Other Steps dialog box opens. The Other Steps dialog box provides instructions on the steps to complete outside the ASDM to make the TLS Proxy fully functional (see Add TLS Proxy Instance Wizard – Other Steps, page 16-11).

Add TLS Proxy Instance Wizard – Other Steps

> ✎
> **Note**   This feature is not supported for the Adaptive Security Appliance version 8.1.2.

The last dialog box of the Add TLS Proxy Instance Wizard specifies the additional steps required to make TLS Proxy fully functional. In particular, you need to perform the following tasks to complete the TLS Proxy configuration:

- Export the local CA certificate or LDC Issuer and install them on the original TLS server.

To export the LDC Issuer, go to Configuration > Firewall > Advanced > Certificate Management > Identity Certificates > Export. See the general operations configuration guide.

- For the TLS Proxy, enable Skinny and SIP inspection between the TLS server and TLS clients. See SIP Inspection, page 10-20 and Skinny (SCCP) Inspection, page 10-32. When you are configuring the TLS Proxy for Presence Federation (which uses CUP), you only enable SIP inspection because the feature supports only the SIP protocol.

- For the TLS Proxy for CUMA, enable MMP inspection.

- When using the internal Certificate Authority of the ASA to sign the LDC Issuer for TLS clients, perform the following:

  – Use the Cisco CTL Client to add the server proxy certificate to the CTL file and install the CTL file on the ASA.

    For information on the Cisco CTL Client, see "Configuring the Cisco CTL Client" in *Cisco Unified CallManager Security Guide*.

    http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/security/5_0_4/secuauth.html

    To install the CTL file on the ASA, go to Configuration > Firewall > Unified Communications > CTL Provider > Add. The Add CTL Provider dialog box opens. For information on using this dialog box to install the CTL file, see Add/Edit CTL Provider, page 16-6.

  – Create a CTL provider instance for connections from the CTL clients. See Add/Edit CTL Provider, page 16-6.

Edit TLS Proxy Instance – Server Configuration

Note   This feature is not supported for the Adaptive Security Appliance version 8.1.2.

The TLS Proxy enables inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and to support the Cisco Unified Communications features on the ASA.

Use the Edit TLS Proxy – Server Configuration tab to edit the server proxy parameters for the original TLS Server—the Cisco Unified Call Manager (CUCM) server, the Cisco Unified Presence Server (CUPS), or the Cisco Unified Mobility Advantage (CUMA) server.

Step 1   Open the Configuration > Firewall > Unified Communications > TLS Proxy pane.

Step 2   To edit a TLS Proxy Instance, click **Edit**.

The Edit TLS Proxy Instance dialog box opens.

Step 3   If necessary, click the **Server Configuration** tab.

Step 4   Specify the server proxy certificate by doing one of the following:

- To add a new certificate, click **Manage**. The Manage Identify Certificates dialog box opens.

  When the Phone Proxy is operating in a mixed-mode CUCM cluster, you must import the CUCM certificate by clicking **Add** in the Manage Identify Certificates dialog box.

- To select an existing certificate, select one from the drop-down list.

When you are configuring the TLS Proxy for the Phone Proxy, select the certificate that has a filename beginning with **_internal_PP_**. When you create the CTL file for the Phone Proxy, the ASA, creates an internal trustpoint used by the Phone Proxy to sign the TFTP files. The trustpoint is named **_internal_PP**_*ctl-instance_filename*.

The server proxy certificate is used to specify the trustpoint to present during the TLS handshake. The trustpoint can be self-signed or enrolled locally with the certificate service on the proxy. For example, for the Phone Proxy, the server proxy certificate is used by the Phone Proxy during the handshake with the IP phones.

Step 5   To install the TLS server certificate in the ASA trust store, so that the ASA can authenticate the TLS server during TLS handshake between the proxy and the TLS server, click **Install TLS Server's Certificate**.

The Manage CA Certificates dialog box opens. Click **Add** to open the Install Certificate dialog box.

When you are configuring the TLS Proxy for the Phone Proxy, click **Install TLS Server's Certificate** and install the Cisco Unified Call Manager (CUCM) certificate so that the proxy can authenticate the IP phones on behalf of the CUCM server.

Step 6   To require the ASA to present a certificate and authenticate the TLS client during TLS handshake, check the Enable client authentication during TLS Proxy handshake check box.

When adding a TLS Proxy Instance for Mobile Advantage (the CUMC client and CUMA server), disable the check box when the client is incapable of sending a client certificate.

Step 7   Click **Apply** to save the changes.

Edit TLS Proxy Instance – Client Configuration

✎
Note   This feature is not supported for the Adaptive Security Appliance version 8.1.2.

The TLS Proxy enables inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and to support the Cisco Unified Communications features on the ASA.

The fields in the Edit TLS Proxy dialog box are identical to the fields displayed when you add a TLS Proxy instance. Use the Edit TLS Proxy – Client Configuration tab to edit the client proxy parameters for the original TLS Client, such as IP phones, CUMA clients, the Cisco Unified Presence Server (CUPS), or the Microsoft OCS server.

Step 1   Open the Configuration > Firewall > Unified Communications > TLS Proxy pane.

Step 2   To edit a TLS Proxy Instance, click **Edit**.

The Edit TLS Proxy Instance dialog box opens.

Step 3   If necessary, click the **Client Configuration** tab.

Step 4   To specify a client proxy certificate to use for the TLS Proxy, perform the following. Select this option when the client proxy certificate is being used between two servers; for example, when configuring the TLS Proxy for Presence Federation, which uses the Cisco Unified Presence Server (CUPS), both the TLS client and TLS server are both servers.

   a.   Check the Specify the proxy certificate for the TLS Client... check box.

   b.   Select a certificate from the drop-down list.

Or

To create a new client proxy certificate, click **Manage**. The Manage Identify Certificates dialog box opens.

> **Note** When you are configuring the TLS Proxy for the Phone Proxy and it is using the mixed security mode for the CUCM cluster, you must configure the LDC Issuer. The LDC Issuer lists the local certificate authority to issue client or server dynamic certificates.

Step 5 To specify an LDC Issuer to use for the TLS Proxy, perform the following. When you select and configure the LDC Issuer option, the ASA acts as the certificate authority and issues certificates to TLS clients.

    a. Click the Specify the internal Certificate Authority to sign the local dynamic certificate for phones... check box.

    b. Click the Certificates radio button and select a self-signed certificate from the drop-down list or click **Manage** to create a new LDC Issuer. The Manage Identify Certificates dialog box opens.

    Or

    Click the Certificate Authority radio button to specify a Certificate Authority (CA) server. When you specify a CA server, it needs to be created and enabled in the ASA. To create and enable the CA server, click **Manage**. The Edit CA Server Settings dialog box opens.

> **Note** To make configuration changes after the local certificate authority has been configured for the first time, disable the local certificate authority.

    c. In the Key-Pair Name field, select a key pair from the drop-list. The list contains the already defined RSA key pair used by client dynamic certificates. To see key pair details, including generation time, usage, modulus size, and key data, click **Show**.

    Or

    To create a new key pair, click **New**. The Add Key Pair dialog box opens.

Step 6 In the Security Algorithms area, specify the available and active algorithms to be announced or matched during the TLS handshake.

- Available Algorithms—Lists the available algorithms to be announced or matched during the TLS handshake: des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, and null-sha1.

  Add—Adds the selected algorithm to the active list.

  Remove—Removes the selected algorithm from the active list.

- Active Algorithms—Lists the active algorithms to be announced or matched during the TLS handshake: des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, and null-sha1. For client proxy (acting as a TLS client to the server), the user-defined algorithms replace the original ones from the hello message for asymmetric encryption method between the two TLS legs. For example, the leg between the proxy and Call Manager may be NULL cipher to offload the Call Manager.

  Move Up—Moves an algorithm up in the list.

  Move Down—Moves an algorithm down in the list.

Step 7 Click **Apply** to save the changes.

TLS Proxy

This feature is supported only for ASA versions 8.0.x prior to 8.0.4 and for version 8.1.

Note   This feature is not supported for the Adaptive Security Appliance versions prior to 8.0.4 and for version 8.1.2.

Use the TLS Proxy option to enable inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco CallManager.

The TLS Proxy pane lets you define and configure Transaction Layer Security Proxy to enable inspection of encrypted traffic.

Fields

- TLS Proxy Name—Lists the TLS Proxy name.
- Server—Lists the trustpoint, which is either self-signed or enrolled with a certificate server.
- Local Dynamic Certificate Issuer—Lists the local certificate authority to issue client or server dynamic certificates.
- Local Dynamic Certificate Key Pair—Lists the RSA key pair used by client or server dynamic certificates.
- Add—Adds a TLS Proxy.
- Edit—Edits a TLS Proxy.
- Delete—Deletes a TLS Proxy.
- Maximum Sessions—Lets you specify the maximum number of TLS Proxy sessions to support.
    - Specify the maximum number of TLS Proxy sessions that the ASA needs to support. By default, ASA supports 300 sessions.—Enables maximum number of sessions option.
    - Maximum number of sessions:—The minimum is 1. The maximum is dependent on the platform. The default is 300.

Add/Edit TLS Proxy

Note   This feature is not supported for the Adaptive Security Appliance versions prior to 8.0.4 and for version 8.1.2.

The Add/Edit TLS Proxy dialog box lets you define the parameters for the TLS Proxy.

Fields

- TLS Proxy Name—Specifies the TLS Proxy name.
- Server Configuration—Specifies the proxy certificate name.
    - Server—Specifies the trustpoint to be presented during the TLS handshake. The trustpoint could be self-signed or enrolled locally with the certificate service on the proxy.
- Client Configuration—Specifies the local dynamic certificate issuer and key pair.
    - Local Dynamic Certificate Issuer—Lists the local certificate authority to issue client or server dynamic certificates.

Certificate Authority Server—Specifies the certificate authority server.

Certificate—Specifies a certificate.

Manage—Configures the local certificate authority. To make configuration changes after it has been configured for the first time, disable the local certificate authority.

– Local Dynamic Certificate Key Pair—Lists the RSA key pair used by client dynamic certificates.

Key-Pair Name—Specifies a defined key pair.

Show—Shows the key pair details, including generation time, usage, modulus size, and key data.

New—Lets you define a new key pair.

- More Options—Specifies the available and active algorithms to be announced or matched during the TLS handshake.

– Available Algorithms—Lists the available algorithms to be announced or matched during the TLS handshake: des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, and null-sha1.

Add—Adds the selected algorithm to the active list.

Remove—Removes the selected algorithm from the active list.

– Active Algorithms—Lists the active algorithms to be announced or matched during the TLS handshake: des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, and null-sha1. For client proxy (acting as a TLS client to the server), the user-defined algorithms replace the original ones from the hello message for asymmetric encryption method between the two TLS legs. For example, the leg between the proxy and CallManager may be NULL cipher to offload the CallManager.

Move Up—Moves an algorithm up in the list.

Move Down—Moves an algorithm down in the list.

Feature History for the TLS Proxy for Encrypted Voice Inspection

Table 16-2 lists the release history for this feature.

Table 16-2   Feature History for Cisco Phone Proxy

Feature Name	Releases	Feature Information
TLS Proxy	8.0(2)	The TLS proxy feature was introduced.

CHAPTER 17

ASA and Cisco Mobility Advantage

This chapter describes how to configure the ASA for Cisco Unified Communications Mobility Advantage Proxy features.

This chapter includes the following sections:

Information about the Cisco Mobility Advantage Proxy Feature

This section contains the following topics:

Cisco Mobility Advantage Proxy Functionality

To support Cisco UMA for the Cisco Mobility Advantage solution, the mobility advantage proxy (implemented as a TLS proxy) includes the following functionality:

- The ability to allow no client authentication during the handshake with clients.
- Allowing an imported PKCS-12 certificate to server as a proxy certificate.

The ASA includes an inspection engine to validate the Cisco UMA Mobile Multiplexing Protocol (MMP).

MMP is a data transport protocol for transmitting data entities between Cisco UMA clients and servers. As shown in Figure 17-1, MMP must be run on top of a connection-oriented protocol (the underlying transport) and is intended to be run on top of a secure transport protocol such as TLS. The Orative Markup Language (OML) protocol is intended to be run on top of MMP for the purposes of data synchronization, as well as the HTTP protocol for uploading and downloading large files.

Figure 17-1   MMP Stack

The TCP/TLS default port is 5443. There are no embedded NAT or secondary connections.

Cisco UMA client and server communications can be proxied via TLS, which decrypts the data, passes it to the inspect MMP module, and re-encrypt the data before forwarding it to the endpoint. The inspect MMP module verifies the integrity of the MMP headers and passes the OML/HTTP to an appropriate handler. The ASA takes the following actions on the MMP headers and data:

- Verifies that client MMP headers are well-formed. Upon detection of a malformed header, the TCP session is terminated.

- Verifies that client to server MMP header lengths are not exceeded. If an MMP header length is exceeded (4096), then the TCP session is terminated.

- Verifies that client to server MMP content lengths are not exceeded. If an entity content length is exceeded (4096), the TCP session is terminated.

Note   4096 is the value currently used in MMP implementations.

Because MMP headers and entities can be split across packets, the ASA buffers data to ensure consistent inspection. The SAPI (stream API) handles data buffering for pending inspection opportunities. MMP header text is treated as case insensitive and a space is present between header text and values. Reclaiming of MMP state is performed by monitoring the state of the TCP connection.

Mobility Advantage Proxy Deployment Scenarios

Figure 17-2 and Figure 17-3 show the two deployment scenarios for the TLS proxy used by the Cisco Mobility Advantage solution. In scenario 1 (the recommended deployment architecture), the ASA functions as both the firewall and TLS proxy. In scenario 2, the ASA functions as the TLS proxy only and works with an existing firewall. In both scenarios, the clients connect from the Internet.

In the scenario 1 deployment, the ASA is between a Cisco UMA client and a Cisco UMA server. The Cisco UMA client is an executable that is downloaded to each smartphone. The Cisco UMA client applications establishes a data connection, which is a TLS connection, to the corporate Cisco UMA server. The ASA intercepts the connections and inspects the data that the client sends to the Cisco UMA server.

Figure 17-2   The TLS proxy for the Cisco Mobility Advantage solution does not support client authentication because the Cisco UMA client cannot present a certificate. **Security Appliance as Firewall with Mobility Advantage Proxy and MMP Inspection**
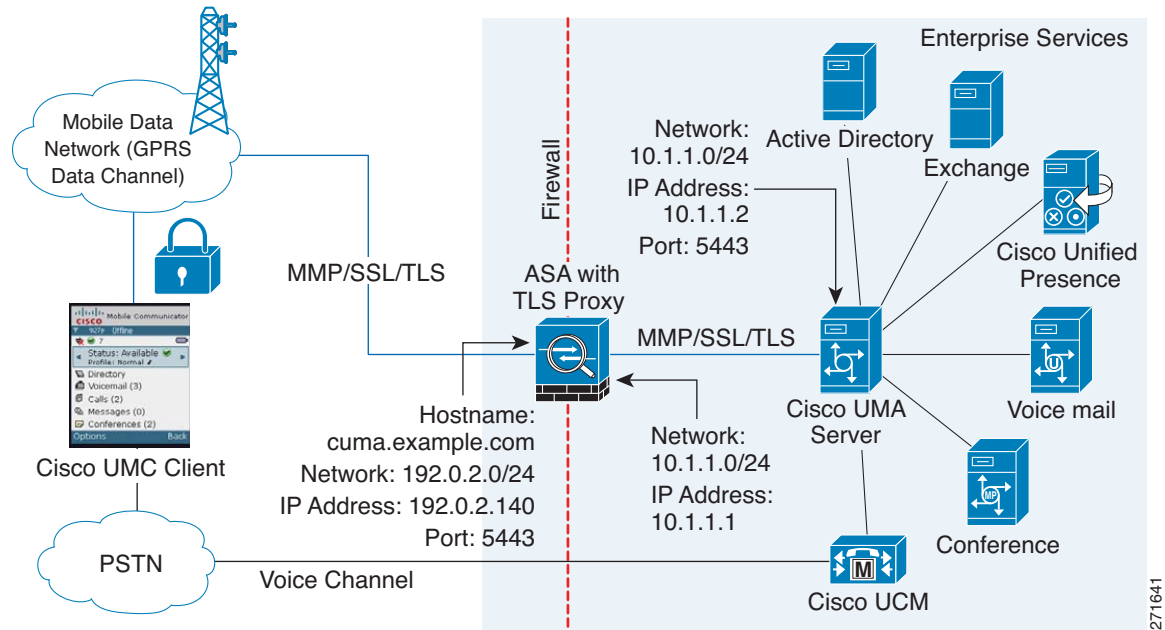
In Figure 17-2, the ASA performs static NAT by translating the Cisco UMA server 10.1.1.2 IP address to 192.0.2.140.

Figure 17-3 shows deployment scenario 2, where the ASA functions as the TLS proxy only and does not function as the corporate firewall. In this scenario, the ASA and the corporate firewall are performing NAT. The corporate firewall will not be able to predict which client from the Internet needs to connect to the corporate Cisco UMA server. Therefore, to support this deployment, you can take the following actions:

- Set up a NAT rule for inbound traffic that translates the destination IP address 192.0.2.41 to 172.16.27.41.

- Set up an interface PAT rule for inbound traffic translating the source IP address of every packet so that the corporate firewall does not need to open up a wildcard pinhole. The Cisco UMA server receives packets with the source IP address 192.0.12.183.

  See Chapter 6, "Network Object NAT (ASA 8.3 and Later)" and Chapter 7, "Twice NAT (ASA 8.3 and Later)" for information.

✎ **Note**   This interface PAT rule converges the Cisco UMA client IP addresses on the outside interface of the ASA into a single IP address on the inside interface by using different source ports. Performing this action is often referred as "outside PAT". "Outside PAT" is not recommended when TLS proxy for Cisco Mobility Advantage is enabled on the same interface of the ASA with phone proxy, Cisco Unified Presence, or any other features involving application inspection. "Outside PAT" is not supported completely by application inspection when embedded address translation is needed.

Figure 17-3   *Cisco UMC/Cisco UMA Architecture – Scenario 2: Security Appliance as Mobility Advantage Proxy Only*
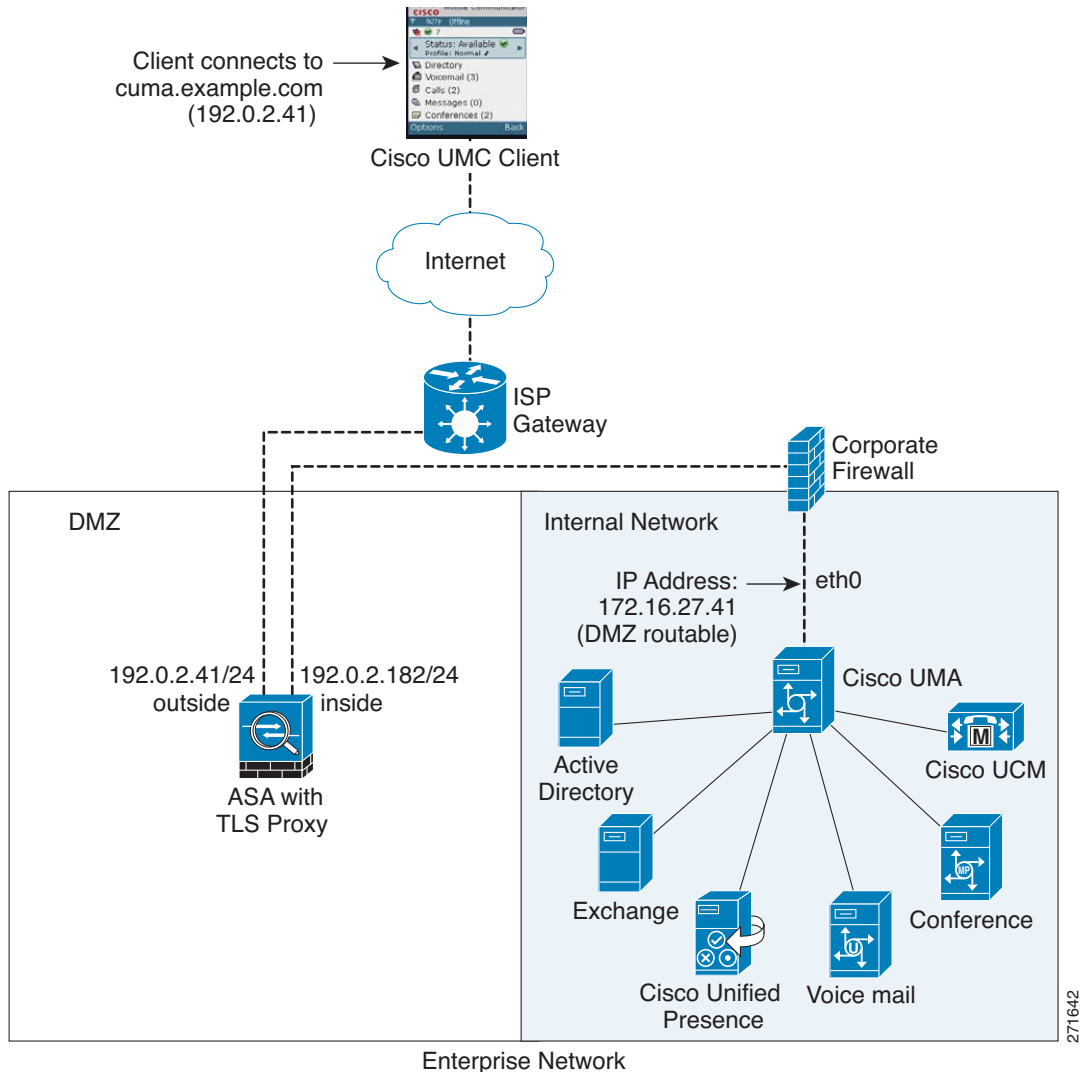
Mobility Advantage Proxy Using NAT/PAT

In both scenarios (Figure 17-2 and Figure 17-3), NAT can be used to hide the private address of the Cisco UMA servers.

In scenario 2 (Figure 17-3), PAT can be used to converge all client traffic into one source IP, so that the firewall does not have to open up a wildcard pinhole for inbound traffic.

Trust Relationships for Cisco UMA Deployments

To establish a trust relationship between the Cisco UMC client and the ASA, the ASA uses the Cisco UMA server certificate and keypair or the ASA obtains a certificate with the Cisco UMA server FQDN (certificate impersonation). Between the ASA and the Cisco UMA server, the ASA and Cisco UMA server use self-signed certificates or certificates issued by a local certificate authority.

Figure 17-4 shows how you can import the Cisco UMA server certificate onto the ASA. When the Cisco UMA server has already enrolled with a third-party CA, you can import the certificate with the private key onto the ASA. Then, the ASA has the full credentials of the Cisco UMA server. When a Cisco UMA client connects to the Cisco UMA server, the ASA intercepts the handshake and uses the Cisco UMA server certificate to perform the handshake with the client. The ASA also performs a handshake with the server.

Figure 17-4   How the Security Appliance Represents Cisco UMA – Private Key Sharing

Figure 17-5 shows another way to establish the trust relationship. Figure 17-5 shows a green field deployment, because each component of the deployment has been newly installed. The ASA enrolls with the third-party CA by using the Cisco UMA server FQDN as if the ASA is the Cisco UMA server. When the Cisco UMA client connects to the ASA, the ASA presents the certificate that has the Cisco UMA server FQDN. The Cisco UMA client believes it is communicating to with the Cisco UMA server.

Figure 17-5   How the Security Appliance Represents Cisco UMA – Certificate Impersonation

A trusted relationship between the ASA and the Cisco UMA server can be established with self-signed certificates. The ASA's identity certificate is exported, and then uploaded on the Cisco UMA server truststore. The Cisco UMA server certificate is downloaded, and then uploaded on the ASA truststore by creating a trustpoint and using the **crypto ca authenticate** command.

Licensing for the Cisco Mobility Advantage Proxy Feature

Note   This feature is not available on No Payload Encryption models.

Model	License Requirement[1]
ASA 5505	Base License and Security Plus License: 2 sessions. *Optional license: 24 sessions.*
ASA 5512-X	Base License or Security Plus License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5515-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5525-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, or 1000 sessions.*
ASA 5545-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, or 2000 sessions.*

Model	License Requirement[1]
ASA 5555-X	Base License: 2 sessions.
	Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.
ASA 5585-X with SSP-10	Base License: 2 sessions.
	Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.
ASA 5585-X with SSP-20, -40, or -60	Base License: 2 sessions.
	Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.
ASASM	Base License: 2 sessions.
	Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.
ASAv with 1 Virtual CPU	Standard and Premium Licenses: 250 sessions.
ASAv with 4 Virtual CPUs	Standard and Premium Licenses: 1000 sessions.

1. The following applications use TLS proxy sessions for their connections. Each TLS proxy session used by these applications (and only these applications) is counted against the UC license limit:
    - Phone Proxy
    - Presence Federation Proxy
    - Encrypted Voice Inspection

    Other applications that use TLS proxy sessions do not count towards the UC limit, for example, Mobility Advantage Proxy (which does not require a license) and IME (which requires a separate IME license).

    Some UC applications might use multiple sessions for a connection. For example, if you configure a phone with a primary and backup Cisco Unified Communications Manager, there are 2 TLS proxy connections, so 2 UC Proxy sessions are used.

    You independently set the TLS proxy limit using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. When you apply a UC license that is higher than the default TLS proxy limit, the ASA automatically sets the TLS proxy limit to match the UC limit. The TLS proxy limit takes precedence over the UC license limit; if you set the TLS proxy limit to be less than the UC license, then you cannot use all of the sessions in your UC license.

    Note: For license part numbers ending in "K8" (for example, licenses under 250 users), TLS proxy sessions are limited to 1000. For license part numbers ending in "K9" (for example, licenses 250 users or larger), the TLS proxy limit depends on the configuration, up to the model limit. K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted.

    Note: If you clear the configuration, then the TLS proxy limit is set to the default for your model; if this default is lower than the UC license limit, then you see an error message to use the to raise the limit again (in ASDM, use the **TLS Proxy** pane). If you use failover and use **File > Save Running Configuration to Standby Unit** on the primary unit to force a configuration synchronization, the **clear configure all** command is generated on the secondary unit automatically, so you may see the warning message on the secondary unit. Because the configuration synchronization restores the TLS proxy limit set on the primary unit, you can ignore the warning.

    You might also use SRTP encryption sessions for your connections:
    - For K8 licenses, SRTP sessions are limited to 250.
    - For K9 licenses, there is not limit.

    Note: Only calls that require encryption/decryption for media are counted towards the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count towards the limit.

Configuring Cisco Mobility Advantage

This section includes the following topic:

Task Flow for Configuring Cisco Mobility Advantage

To configure for the ASA to perform TLS proxy and MMP inspection as shown in Figure 17-2 and Figure 17-3, perform the following tasks.

It is assumed that self-signed certificates are used between the ASA and the Cisco UMA server.

To configure the Cisco Mobility Advantage Proxy by using ASDM, choose Wizards > Unified Communications Wizard from the menu. The Unified Communications Wizard opens. From the first page, select the Cisco Mobility Advantage Proxy option under the Remote Access section.

The wizard automatically creates the necessary TLS proxy, then guides you through creating the Unified Presence Proxy instance, importing and installing the required certificates, and finally enables the MMP inspection for the Mobility Advantage traffic automatically.

The wizard guides you through four steps to create the Mobility Advantage Proxy:

Step 1 Select the Mobility Advantage Proxy option.

Step 2 Specify setting to define the proxy topology, such the IP address of the Mobility Advantage server.

Step 3 Configure the server-side certificate management, namely the certificates that are exchanged between the local Mobility Advantage server and the ASA.

Step 4 Configure the client-side certificate management, namely the certificates that are exchanged between the Unified Mobile Communicator and the ASA

The wizard completes by displaying a summary of the configuration created for Mobility Advantage Proxy. See Chapter 14, "Cisco Unified Communication Wizard" for more information.

Feature History for Cisco Mobility Advantage

Table 17-1 lists the release history for this feature.

Table 17-1   Feature History for Cisco Phone Proxy

Feature Name	Releases	Feature Information
Cisco Mobility Advantage Proxy	8.0(4)	The Cisco Mobility Advantage Proxy feature was introduced.
Cisco Mobility Advantage Proxy	8.3(1)	The Unified Communications Wizard was added to ASDM. By using the wizard, you can configure the Cisco Mobility Advantage Proxy.

CHAPTER **18**

ASA and Cisco Unified Presence

This chapter describes how to configure the ASA for Cisco Unified Presence.

This chapter includes the following sections:

- Information About Cisco Unified Presence, page 18-1
- Licensing for Cisco Unified Presence, page 18-7
- Configuring Cisco Unified Presence Proxy for SIP Federation, page 18-8
- Feature History for Cisco Unified Presence, page 18-9

Information About Cisco Unified Presence

This section includes the following topics:

- Architecture for Cisco Unified Presence for SIP Federation Deployments, page 18-1
- Trust Relationship in the Presence Federation, page 18-4
- Security Certificate Exchange Between Cisco UP and the Security Appliance, page 18-5
- XMPP Federation Deployments, page 18-5
- Configuration Requirements for XMPP Federation, page 18-6

Architecture for Cisco Unified Presence for SIP Federation Deployments

Figure 18-1 depicts a Cisco Unified Presence/LCS Federation scenario with the ASA as the presence federation proxy (implemented as a TLS proxy). The two entities with a TLS connection are the "Routing Proxy" (a dedicated Cisco UP) in Enterprise X and the Microsoft Access Proxy in Enterprise Y. However, the deployment is not limited to this scenario. Any Cisco UP or Cisco UP cluster could be deployed on the left side of the ASA; the remote entity could be any server (an LCS, an OCS, or another Cisco UP).

The following architecture is generic for two servers using SIP (or other ASA inspected protocols) with a TLS connection.

Entity X: Cisco UP/Routing Proxy in Enterprise X

Entity Y: Microsoft Access Proxy/Edge server for LCS/OCS in Enterprise Y

Figure 18-1   Typical Cisco Unified Presence/LCS Federation Scenario

In the above architecture, the ASA functions as a firewall, NAT, and TLS proxy, which is the recommended architecture. However, the ASA can also function as NAT and the TLS proxy alone, working with an existing firewall.

Either server can initiate the TLS handshake (unlike IP Telephony or Cisco Unified Mobility, where only the clients initiate the TLS handshake). There are by-directional TLS proxy rules and configuration. Each enterprise can have an ASA as the TLS proxy.

In Figure 18-1, NAT or PAT can be used to hide the private address of Entity X. In this situation, static NAT or PAT must be configured for foreign server (Entity Y) initiated connections or the TLS handshake (inbound). Typically, the public port should be 5061. The following static PAT command is required for the Cisco UP that accepts inbound connections:

```
hostname(config)# object network obj-10.0.0.2-01
hostname(config-network-object)# host 10.0.0.2
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service tcp 5061
5061
```

The following static PAT must be configured for each Cisco UP that could initiate a connection (by sending SIP SUBSCRIBE) to the foreign server.

For Cisco UP with the address 10.0.0.2, enter the following command:

```
hostname(config)# object network obj-10.0.0.2-02
hostname(config-network-object)# host 10.0.0.2
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service tcp 5062
5062
hostname(config)# object network obj-10.0.0.2-03
hostname(config-network-object)# host 10.0.0.2
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service udp 5070
5070
hostname(config)# object network obj-10.0.0.2-04
hostname(config-network-object)# host 10.0.0.2
```

```
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service tcp 5060
5060
```

For another Cisco UP with the address 10.0.0.3, you must use a different set of PAT ports, such as 45062 or 45070:

```
hostname(config)# object network obj-10.0.0.3-01
hostname(config-network-object)# host 10.0.0.3
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service tcp 5061
45061
hostname(config)# object network obj-10.0.0.3-02
hostname(config-network-object)# host 10.0.0.3
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service tcp 5062
45062
hostname(config)# object network obj-10.0.0.3-03
hostname(config-network-object)# host 10.0.0.3
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service udp 5070
5070
hostname(config)# object network obj-10.0.0.2-03
hostname(config-network-object)# host 10.0.0.2
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service tcp 5070
45070
hostname(config)# object network obj-10.0.0.3-04
hostname(config-network-object)# host 10.0.0.3
hostname(config-network-object)# nat (inside,outside) static 192.0.2.1 service tcp 5060
45060
```

Dynamic NAT or PAT can be used for the rest of the outbound connections or the TLS handshake. The ASA SIP inspection engine takes care of the necessary translation (fixup).

```
hostname(config)# object network obj-0.0.0.0-01
hostname(config-network-object)# subnet 0.0.0.0 0.0.0.0
hostname(config-network-object)# nat (inside,outside) dynamic 192.0.2.1
```

Figure 18-2 illustrates an abstracted scenario with Entity X connected to Entity Y through the presence federation proxy on the ASA. The proxy is in the same administrative domain as Entity X. Entity Y could have another ASA as the proxy but this is omitted for simplicity.

Figure 18-2   Abstracted Presence Federation Proxy Scenario between Two Server Entities

For the Entity X domain name to be resolved correctly when the ASA holds its credential, the ASA could be configured to perform NAT for Entity X, and the domain name is resolved as the Entity X public address for which the ASA provides proxy service.

For further information about configuring Cisco Unified Presence Federation for SIP Federation, see the Integration Guide for Configuring Cisco Unified Presence for Interdomain Federation.:

http://www.cisco.com/en/US/products/ps6837/products_installation_and_configuration_guides_list.html

Trust Relationship in the Presence Federation

Within an enterprise, setting up a trust relationship is achievable by using self-signed certificates or you can set it up on an internal CA.

Establishing a trust relationship cross enterprises or across administrative domains is key for federation. Cross enterprises you must use a trusted third-party CA (such as, VeriSign). The ASA obtains a certificate with the FQDN of the Cisco UP (certificate impersonation).

For the TLS handshake, the two entities could validate the peer certificate via a certificate chain to trusted third-party certificate authorities. Both entities enroll with the CAs. The ASA as the TLS proxy must be trusted by both entities. The ASA is always associated with one of the enterprises. Within that enterprise (Enterprise X in Figure 18-1), the entity and the ASA could authenticate each other via a local CA, or by using self-signed certificates.

To establish a trusted relationship between the ASA and the remote entity (Entity Y), the ASA can enroll with the CA on behalf of Entity X (Cisco UP). In the enrollment request, the Entity X identity (domain name) is used.

Figure 18-3 shows the way to establish the trust relationship. The ASA enrolls with the third party CA by using the Cisco UP FQDN as if the ASA is the Cisco UP.

Figure 18-3   How the Security Appliance Represents Cisco Unified Presence – Certificate Impersonate
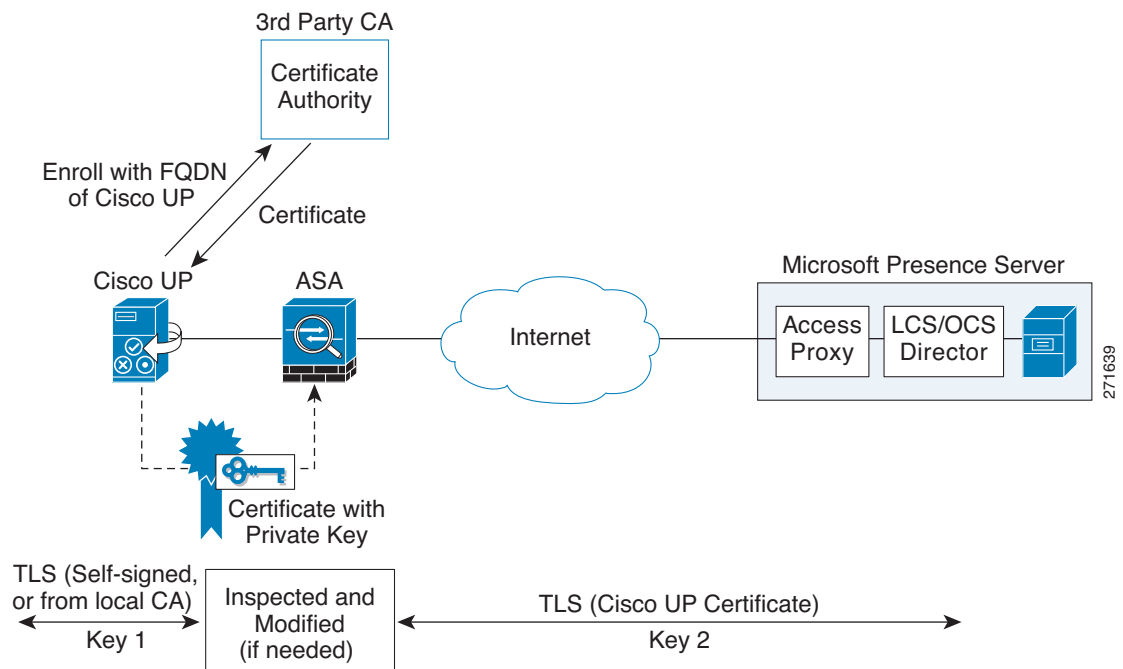
Security Certificate Exchange Between Cisco UP and the Security Appliance

You need to generate the keypair for the certificate (such as `cup_proxy_key`) used by the ASA, and configure a trustpoint to identify the self-signed certificate sent by the ASA to Cisco UP (such as `cup_proxy`) in the TLS handshake.

For the ASA to trust the Cisco UP certificate, you need to create a trustpoint to identify the certificate from the Cisco UP (such as `cert_from_cup`), and specify the enrollment type as terminal to indicate that you will paste the certificate received from the Cisco UP into the terminal.

XMPP Federation Deployments

Figure 18-4 provides an example of an XMPP federated network between Cisco Unified Presence enterprise deployment and an IBM Sametime enterprise deployment. TLS is optional for XMPP federation. ASA acts only as a firewall for XMPP federation; it does not provide TLS proxy functionality or PAT for XMPP federation.

Figure 18-4   Basic XMPP Federated Network between Cisco Unified Presence and IBM Sametime
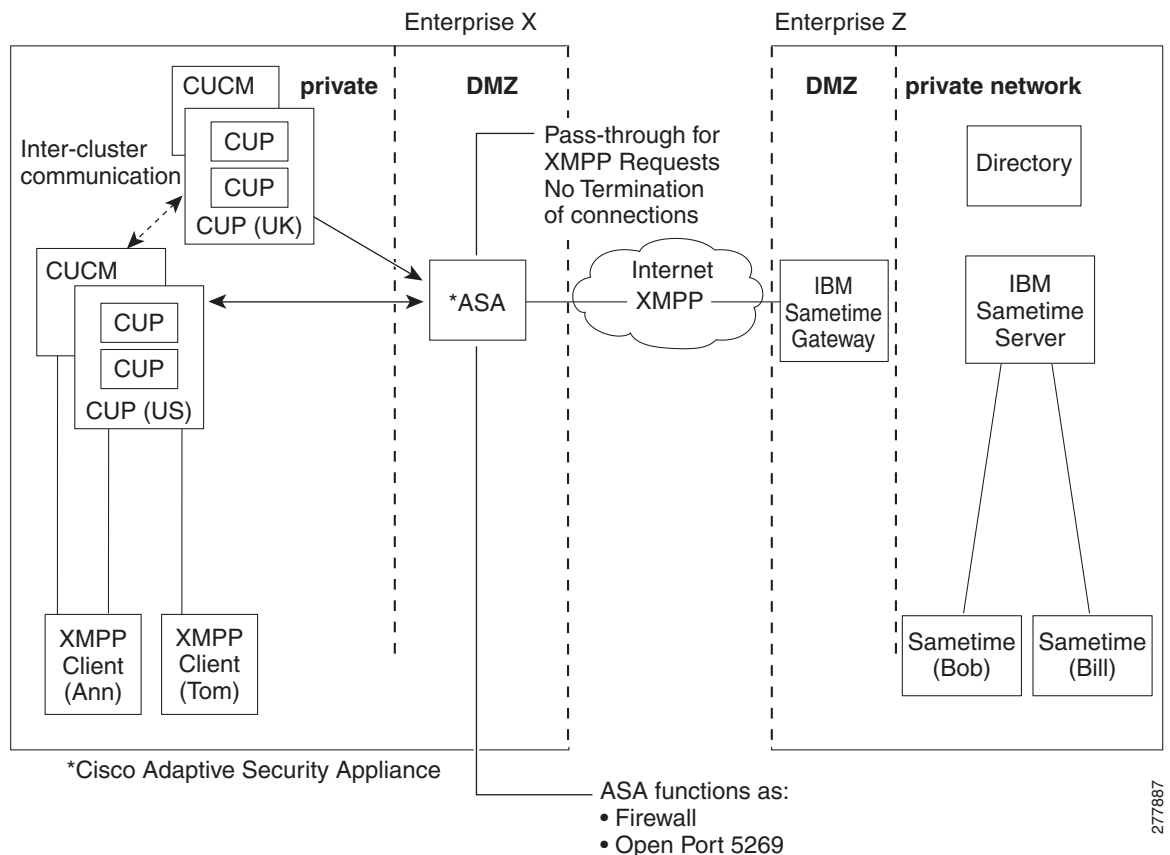
There are two DNS servers within the internal Cisco Unified Presence enterprise deployment. One DNS server hosts the Cisco Unified Presence private address. The other DNS server hosts the Cisco Unified Presence public address and a DNS SRV records for SIP federation (_sipfederationtle), and XMPP federation (_xmpp-server) with Cisco Unified Presence. The DNS server that hosts the Cisco Unified Presence public address is located in the local DMZ.

For further information about configuring Cisco Unified Presence Federation for XMPP Federation, see the *Integration Guide for Configuring Cisco Unified Presence Release 8.0 for Interdomain Federation*:

http://www.cisco.com/en/US/products/ps6837/products_installation_and_configuration_guides_list.html

Configuration Requirements for XMPP Federation

For XMPP Federation, ASA acts as a firewall only. You must open port 5269 for both incoming and outgoing XMPP federated traffic on ASA.

These are sample ACLs to open port 5269 on ASA.

Allow traffic from any address to any address on port 5269:

```
access-list ALLOW-ALL extended permit tcp any any eq 5269
```

Allow traffic from any address to any single node on port 5269:

```
access-list ALLOW-ALL extended permit tcp any host <private cup IP address> eq 5269
```

If you do not configure the ACL above, and you publish additional XMPP federation nodes in DNS, you must configure access to each of these nodes, for example:

```
object network obj_host_<private cup ip address>
#host <private cup ip address>
object network obj_host_<private cup2 ip address>
#host <private cup2 ip address>
object network obj_host_<public cup ip address>
#host <public cup ip address>
....
```

Configure the following NAT commands:

```
nat (inside,outside) source static obj_host_<private cup1 IP> obj_host_<public cup IP>
service
obj_udp_source_eq_5269 obj_udp_source_eq_5269
nat (inside,outside) source static obj_host_<private cup1 IP> obj_host_<public cup IP>
service
obj_tcp_source_eq_5269 obj_tcp_source_eq_5269
```

If you publish a single public IP address in DNS, and use arbitrary ports, configure the following:

(This example is for two additional XMPP federation nodes)

```
nat (inside,outside) source static obj_host_<private cup2 ip> obj_host_<public cup IP>
service
obj_udp_source_eq_5269 obj_udp_source_eq_25269
nat (inside,outside) source static obj_host_<private cup2 ip> obj_host_<public cup IP>
service
obj_tcp_source_eq_5269 obj_tcp_source_eq_25269

nat (inside,outside) source static obj_host_<private cup3 ip> obj_host_<public cup IP>
service
obj_udp_source_eq_5269 obj_udp_source_eq_35269
nat (inside,outside) source static obj_host_<private cup3 ip> obj_host_<public cup IP>
service
obj_tcp_source_eq_5269 obj_tcp_source_eq_35269
```

If you publish multiple public IP addresses in DNS all using port 5269, configure the following:

(This example is for two additional XMPP federation nodes)

```
nat (inside,outside) source static obj_host_<private cup2 ip> obj_host_<public cup2 IP>
service
obj_udp_source_eq_5269 obj_udp_source_eq_5269
nat (inside,outside) source static obj_host_<private cup2 ip> obj_host_<public cup2 IP>
service
obj_tcp_source_eq_5269 obj_tcp_source_eq_5269

nat (inside,outside) source static obj_host_<private cup3 ip> obj_host_<public cup3 IP>
service
obj_udp_source_eq_5269 obj_udp_source_eq_5269
nat (inside,outside) source static obj_host_<private cup3 ip> obj_host_<public cup IP>
service
obj_tcp_source_eq_5269 obj_tcp_source_eq_5269
```

Licensing for Cisco Unified Presence

The Cisco Unified Presence feature supported by the ASA require a Unified Communications Proxy license.

The following table shows the Unified Communications Proxy license details by platform:

Note This feature is not available on No Payload Encryption models.

Model	License Requirement[1]
ASA 5505	Base License and Security Plus License: 2 sessions. *Optional license: 24 sessions.*
ASA 5512-X	Base License or Security Plus License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5515-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, or 500 sessions.*
ASA 5525-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, or 1000 sessions.*
ASA 5545-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, or 2000 sessions.*
ASA 5555-X	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5585-X with SSP-10	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, or 3000 sessions.*
ASA 5585-X with SSP-20, -40, or -60	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*
ASASM	Base License: 2 sessions. *Optional licenses: 24, 50, 100, 250, 500, 750, 1000, 2000, 3000, 5000, or 10,000 sessions.*

Model	License Requirement[1]
ASAv with 1 Virtual CPU	Standard and Premium Licenses: 250 sessions.
ASAv with 4 Virtual CPUs	Standard and Premium Licenses: 1000 sessions.

1. The following applications use TLS proxy sessions for their connections. Each TLS proxy session used by these applications (and only these applications) is counted against the UC license limit:
    - Phone Proxy
    - Presence Federation Proxy
    - Encrypted Voice Inspection

    Other applications that use TLS proxy sessions do not count towards the UC limit, for example, Mobility Advantage Proxy (which does not require a license) and IME (which requires a separate IME license).

    Some UC applications might use multiple sessions for a connection. For example, if you configure a phone with a primary and backup Cisco Unified Communications Manager, there are 2 TLS proxy connections, so 2 UC Proxy sessions are used.

    You independently set the TLS proxy limit using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. When you apply a UC license that is higher than the default TLS proxy limit, the ASA automatically sets the TLS proxy limit to match the UC limit. The TLS proxy limit takes precedence over the UC license limit; if you set the TLS proxy limit to be less than the UC license, then you cannot use all of the sessions in your UC license.

    Note: For license part numbers ending in "K8" (for example, licenses under 250 users), TLS proxy sessions are limited to 1000. For license part numbers ending in "K9" (for example, licenses 250 users or larger), the TLS proxy limit depends on the configuration, up to the model limit. K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted.

    Note: If you clear the configuration, then the TLS proxy limit is set to the default for your model; if this default is lower than the UC license limit, then you see an error message to use the to raise the limit again (in ASDM, use the **TLS Proxy** pane). If you use failover and use **File > Save Running Configuration to Standby Unit** on the primary unit to force a configuration synchronization, the **clear configure all** command is generated on the secondary unit automatically, so you may see the warning message on the secondary unit. Because the configuration synchronization restores the TLS proxy limit set on the primary unit, you can ignore the warning.

    You might also use SRTP encryption sessions for your connections:
    - For K8 licenses, SRTP sessions are limited to 250.
    - For K9 licenses, there is not limit.

    Note: Only calls that require encryption/decryption for media are counted towards the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count towards the limit.

Configuring Cisco Unified Presence Proxy for SIP Federation

This section contains the following topic:

-

Task Flow for Configuring Cisco Unified Presence Federation Proxy for SIP Federation

To configure a Cisco Unified Presence/LCS Federation scenario with the ASA as the TLS proxy where there is a single Cisco UP that is in the local domain and self-signed certificates are used between the Cisco UP and the ASA (like the scenario shown in Figure 18-1), perform the following tasks.

To configure the Cisco Unified Presence proxy by using ASDM, choose Wizards > Unified Communications Wizard from the menu. The Unified Communications Wizard opens. From the first page, select the Cisco Unified Presence Proxy option under the Business-to-Business section.

The wizard automatically creates the necessary TLS proxy, then guides you through creating the Unified Presence Proxy instance, importing and installing the required certificates, and finally enables the SIP and SCCP inspection for the Presence Federation traffic automatically.

The wizard guides you through four steps to create the Presence Federation Proxy:

Step 1   Select the Presence Federation Proxy option.

Step 2   Specify setting to define the proxy topology, such the IP address of the Presence Federation server.

Step 3   Configure the local-side certificate management, namely the certificates that are exchanged between the local Unified Presence Federation server and the ASA.

Step 4   Configure the remote-side certificate management, namely the certificates that are exchanged between the remote server and the ASA

The wizard completes by displaying a summary of the configuration created for Presence Federation. See the Unified Communications Wizard section in this documentation for more information.

Feature History for Cisco Unified Presence

Table 18-1 lists the release history for this feature.

Table 18-1   *Feature History for Cisco Unified Presence*

Feature Name	Releases	Feature Information
Cisco Presence Federation Proxy	8.0(4)	The Cisco Unified Presence proxy feature was introduced.
Cisco Presence Federation Proxy	8.3(1)	The Unified Communications Wizard was added to ASDM. By using the wizard, you can configure the Cisco Presence Federation Proxy. Support for XMPP Federation was introduced.

CHAPTER **19**

ASA and Cisco Intercompany Media Engine Proxy

This chapter describes how to configure the ASA for Cisco Intercompany Media Engine Proxy.

This chapter includes the following sections:

Information About Cisco Intercompany Media Engine Proxy

This section includes the following topics:

Features of Cisco Intercompany Media Engine Proxy

Cisco Intercompany Media Engine enables companies to interconnect on-demand, over the Internet with advanced features made available by VoIP technologies. Cisco Intercompany Media Engine allows for business-to-business federation between Cisco Unified Communications Manager clusters in different enterprises by utilizing peer-to-peer, security, and SIP protocols to create dynamic SIP trunks between businesses. A collection of enterprises work together to end up looking like one large business with inter-cluster trunks between them.

The adaptive security appliance applies its existing TLS proxy, SIP Application Layer Gateway (ALG), and SIP verification features to the functioning of Cisco Intercompany Media Engine.

Cisco Intercompany Media Engine has the following key features:

- Works with existing phone numbers: Cisco Intercompany Media Engine works with the phone numbers an enterprise currently has and does not require an enterprise to learn new numbers or change providers to use Cisco Intercompany Media Engine.

- Works with existing IP phones: Cisco Intercompany Media Engine works with the existing IP phones within an enterprise. However, the feature set in business-to-business calls is limited to the capabilities of the IP phones.

- Does not require purchasing new services: Cisco Intercompany Media Engine does not require any new services from any service providers. Customers continue to use the PSTN connectivity they have and the Internet connectivity they have today. Cisco Intercompany Media Engine gradually moves calls off the PSTN and onto the Internet.

- Provides a full Cisco Unified Communications experience: Because Cisco Intercompany Media Engine creates inter-cluster SIP trunks between enterprises, any Unified Communication features that work over the SIP trunk and only require a SIP trunk work with the Cisco Intercompany Media Engine, thus providing a Unified Communication experience across enterprises.

- Works on the Internet: Cisco Intercompany Media Engine was designed to work on the Internet. It can also work on managed extranets.

- Provides worldwide reach: Cisco Intercompany Media Engine can connect to any enterprise anywhere in the world, as long as the enterprise is running Cisco Intercompany Media Engine technology. There are no regional limitations. This is because Cisco Intercompany Media Engine utilizes two networks that both have worldwide reach—the Internet and the PSTN.

- Allows for unlimited scale: Cisco Intercompany Media Engine can work with any number of enterprises.

- Is self-learning: The system is primarily self-learning. Customers do not have to enter information about other businesses: no phone prefixes, no IP address, no ports, no domain names, nor certificates. Customers need to configure information about their own networks, and provide policy information if they want to limit the scope of Cisco Intercompany Media Engine.

- Is secure: Cisco Intercompany Media Engine is secure, utilizing a large number of different technologies to accomplish this security.

- Includes anti-spam: Cisco Intercompany Media Engine prevents people from setting up software on the Internet that spams enterprises with phone calls. It provides an extremely high barrier to entry.

- Provides for QoS management: Cisco Intercompany Media Engine provides features that help customers manage the QoS on the Internet, such as the ability to monitor QoS of the RTP traffic in real-time and fallback to PSTN automatically if problems arise.

How the UC-IME Works with the PSTN and the Internet

The Cisco Intercompany Media Engine utilizes two networks that both have worldwide reach—the Internet and the PSTN. Customers continue to use the PSTN connectivity they have. The Cisco Intercompany Media Engine gradually moves calls off the PSTN and onto the Internet. However, if QoS problems arise, the Cisco Intercompany Media Engine Proxy monitors QoS of the RTP traffic in real-time and fallbacks to PSTN automatically.

The Cisco Intercompany Media Engine uses information from PSTN calls to validate that the terminating side owns the number that the originated side had called. After the PSTN call terminates, the enterprises involved in the call send information about the call to their Cisco IME server. The Cisco IME server on the originating side validates the call. Figure 19-1 shows the initial call flow through the PSTN.

On successful verification, the terminating side creates a ticket that grants permission to the call originator to make a Cisco IME call to a specific number. See Tickets and Passwords, page 19-3 for information.

Figure 19-1   Interaction of the UC-IME Proxy with the PSTN

Tickets and Passwords

Cisco Intercompany Media Engine utilizes tickets and passwords to provide enterprise verification. Verification through the creation of tickets ensures an enterprise is not subject to denial-of-service (DOS) attacks from the Internet or endless VoIP spam calls. Ticket verification prevents spam and DOS attacks because it introduces a cost to the VoIP caller; namely, the cost of a PSTN call. A malicious user cannot set up just an open source asterisk PBX on the Internet and begin launching SIP calls into an enterprise running Cisco Intercompany Media Engine. Having the Cisco Intercompany Media Engine Proxy verify tickets allows incoming calls from a particular enterprise to a particular number only when that particular enterprise has previously called that phone number on the PSTN.

To send a spam VoIP call to every phone within an enterprise, an organization would have to purchase the Cisco Intercompany Media Engine and Cisco Unified Communications Manager and have called each phone number within the enterprise over the PSTN and completed each call successfully. Only then can it launch a VoIP call to each number.

The Cisco Intercompany Media Engine server creates tickets and the ASA validates them. The ASA and Cisco Intercompany Media Engine server share a password that is configured so that the ASA detects the ticket was created by a trusted Cisco Intercompany Media Engine server. The ticket contains information that indicates that the enterprise is authorized to call specific phone numbers at the target enterprise. See Figure 19-2 for the ticket verification process and how it operates between the originating and terminating-call enterprises.

Note   Because the initial calls are over the PSTN, they are subject to any national regulations regarding telemarketing calling. For example, within the United States, they would be subject to the national do-not-call registry.

Figure 19-2   Ticket Verification Process with Cisco Intercompany Media Engine
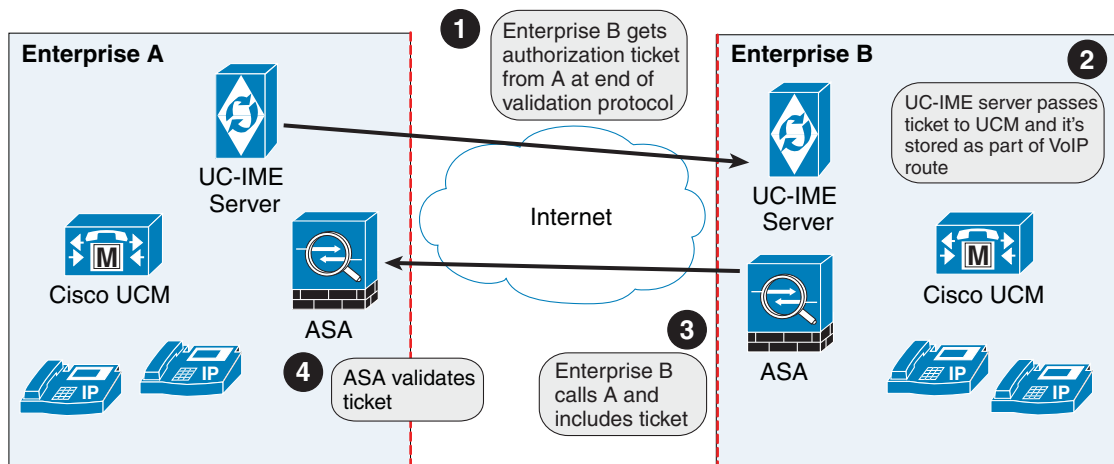
As illustrated in Figure 19-2. Enterprise B makes a PSTN call to enterprise A. That call completes successfully. Later, Enterprise B Cisco Intercompany Media Engine server initiates validation procedures with Enterprise A. These validation procedures succeed. During the validation handshake, Enterprise B sends Enterprise A its domain name. Enterprise A verifies that this domain name is not on the blacklisted set of domains. Assuming it is not, Enterprise A creates a ticket.

Subsequently, someone in Enterprise B calls that number again. That call setup message from Enterprise B to Enterprise A includes the ticket in the X-Cisco-UC-IME-Ticket header field in the SIP INVITE message. This message arrives at the Enterprise A ASA. The ASA verifies the signature and computes several checks on the ticket to make sure it is valid. If the ticket is valid, the ASA forwards the request to Cisco UCM (including the ticket). Because the ASA drops requests that lack a valid ticket, unauthorized calls are never received by Cisco UCM.

The ticket password is a 128 bit random key, which can be thought of as a shared password between the adaptive security appliance and the Cisco Intercompany Media Engine server. This password is generated by the Cisco Intercompany Media Engine server and is used by a Cisco Intercompany Media Engine SIP trunk to generate a ticket to allow a call to be made between Cisco Intercompany Media Engine SIP trunks. A ticket is a signed object that contains a number of fields that grant permission to the calling domain to make a Cisco Intercompany Media Engine call to a specific number. The ticket is signed by the ticket password.

The Cisco Intercompany Media Engine also required that you configure an epoch for the password. The epoch contains an integer that updates each time that the password is changed. When the proxy is configured the first time and a password entered for the first time, enter 1 for the epoch integer. Each time you change the password, increment the epoch to indicate the new password. You must increment the epoch value each time your change the password.

Typically, you increment the epoch sequentially; however, the ASA allows you to choose any value when you update the epoch. If you change the epoch value, the tickets in use at remote enterprises become invalid. The incoming calls from the remote enterprises fallback to the PSTN until the terminating enterprise reissues tickets with the new epoch value and password.

The epoch and password that you configure on the ASA must match the epoch and password configured on the Cisco Intercompany Media Engine server. If you change the password or epoch on the ASA, you must update them on the Cisco Intercompany Media Engine server. See the Cisco Intercompany Media Engine server documentation for information.

Call Fallback to the PSTN

Cisco Intercompany Media Engine provides features that manage the QoS on the Internet, such as the ability to monitor QoS of the RTP traffic in real-time and fallback to PSTN automatically if problems arise. Call fallback from Internet VoIP calls to the public switched telephone network (PSTN) can occur for two reasons changes in connection quality and signal failure for the Cisco Intercompany Media Engine.

Internet connections can vary wildly in their quality and vary over time. Therefore, even if a call is sent over VoIP because the quality of the connection was good, the connection quality might worsen mid-call. To ensure an overall good experience for the end user, Cisco Intercompany Media Engine attempts to perform a mid-call fallback.

Performing a mid-call fallback requires the adaptive security appliance to monitor the RTP packets coming from the Internet and send information into an RTP Monitoring Algorithm (RMA) API, which will indicates to the adaptive security appliance whether fallback is required. If fallback is required, the adaptive security appliance sends a REFER message to Cisco UCM to tell it that it needs to fallback the call to PSTN.

The TLS signaling connections from the Cisco UCM are terminated on the adaptive security appliance and a TCP or TLS connection is initiated to the Cisco UCM. SRTP (media) sent from external IP phones to the internal network IP phone via the adaptive security appliance is converted to RTP. The adaptive security appliance inserts itself into the media path by modifying the SIP signaling messages that are sent over the SIP trunk between Cisco UCMs. TLS (signaling) and SRTP are always terminated on the adaptive security appliance.

If signaling problems occur, the call falls back to the PSTN; however, the Cisco UCM initiates the PSTN fall back and the adaptive security appliance does not send REFER message.

Architecture and Deployment Scenarios for Cisco Intercompany Media Engine

This section includes the following topics:

Architecture

Within the enterprise, Cisco Intercompany Media Engine is deployed with the following components for the following purposes:

- The adaptive security appliance—Enabled with the Cisco Intercompany Media Engine Proxy, provides perimeter security functions and inspects SIP signaling between SIP trunks.

- Cisco Intercompany Media Engine (UC-IME) server— Located in the DMZ, provides an automated provisioning service by learning new VoIP routes to particular phone numbers, and recording those routes in Cisco UCM. The Cisco Intercompany Media Engine server does not perform call control.

- Cisco Unified Communications Manager (Cisco UCM)—Responsible for call control and processing. Cisco UCM connects to the Cisco Intercompany Media Engine server by using the Access Protocol to publish and exchange updates. The architecture can consist of a single Cisco UCM or a Cisco UCM cluster within the enterprise.

- Cisco Intercompany Media Engine (UC-IME) Bootstrap server—Provides a certificate required admission onto the public peer-to-peer network for Cisco Intercompany Media Engine.

Figure 19-3 illustrates the components of the Cisco Intercompany Media Engine in a basic deployment.

Figure 19-3   Cisco Intercompany Media Engine Architecture in a Basic Deployment
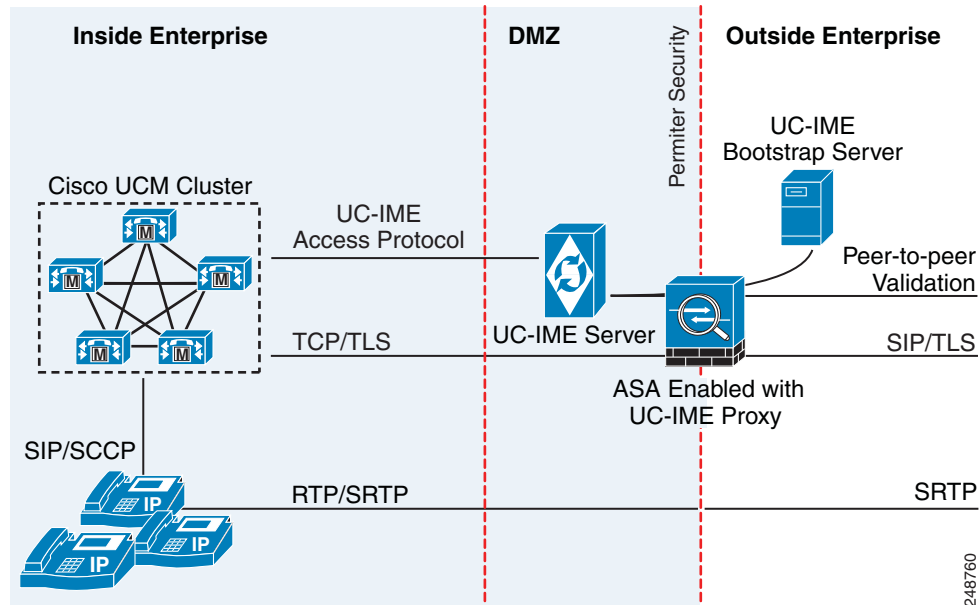
Basic Deployment

In a basic deployment, the Cisco Intercompany Media Engine Proxy sits in-line with the Internet firewall such that all Internet traffic traverses the adaptive security appliance. In this deployment, a single Cisco UCM or a Cisco UCM cluster is centrally deployed within the enterprise, along with a Cisco Intercompany Media Engine server (and perhaps a backup).

As shown in Figure 19-4, the adaptive security appliance sits on the edge of the enterprise and inspects SIP signaling by creating dynamic SIP trunks between enterprises.

Figure 19-4   Basic Deployment Scenario
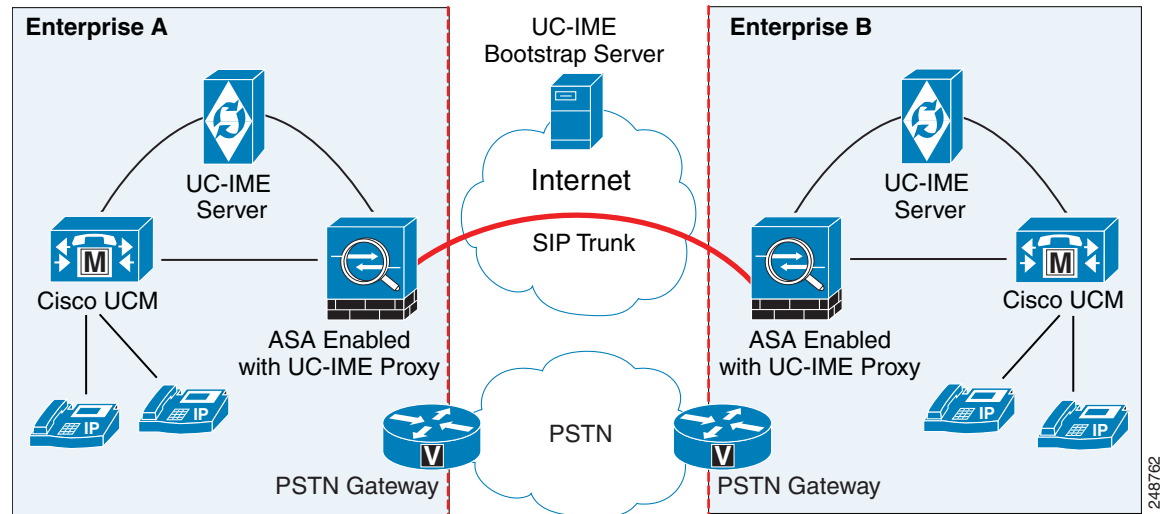
Off Path Deployment

In an off path deployment, inbound and outbound Cisco Intercompany Media Engine calls pass through an adaptive security appliance enabled with the Cisco Intercompany Media Engine Proxy. The adaptive security appliance is located in the DMZ and is configured to support only the Cisco Intercompany Media Engine traffic (SIP signaling and RTP traffic). Normal Internet facing traffic does not flow through this adaptive security appliance.

For all inbound calls, the signaling is directed to the adaptive security appliance because destined Cisco UCMs are configured with the global IP address on the adaptive security appliance. For outbound calls, the called party could be any IP address on the Internet; therefore, the adaptive security appliance is configured with a mapping service that dynamically provides an internal IP address on the adaptive security appliance for each global IP address of the called party on the Internet.

Cisco UCM sends all outbound calls directly to the mapped internal IP address on the adaptive security appliance instead of the global IP address of the called party on the Internet. The adaptive security appliance then forwards the calls to the global IP address of the called party.

Figure 19-5 illustrates the architecture of the Cisco Intercompany Media Engine in an off path deployment.

Figure 19-5   *Off Path Deployment of the Adaptive Security Appliance*

Licensing for Cisco Intercompany Media Engine

The Cisco Intercompany Media Engine feature supported by the ASA require a Unified Communications Proxy license.

The following table shows the details of the Unified Communications Proxy license:

Note   This feature is not available on No Payload Encryption models.

Model	License Requirement
All models	Intercompany Media Engine license.
	When you enable the Intercompany Media Engine (IME) license, you can use TLS proxy sessions up to the configured TLS proxy limit. If you also have a Unified Communications (UC) license installed that is higher than the default TLS proxy limit, then the ASA sets the limit to be the UC license limit plus an additional number of sessions depending on your model. You can manually configure the TLS proxy limit using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. If you also install the UC license, then the TLS proxy sessions available for UC are also available for IME sessions. For example, if the configured limit is 1000 TLS proxy sessions, and you purchase a 750-session UC license, then the first 250 IME sessions do not affect the sessions available for UC. If you need more than 250 sessions for IME, then the remaining 750 sessions of the platform limit are used on a first-come, first-served basis by UC and IME.
	• For a license part number ending in "K8", TLS proxy sessions are limited to 1000.
	• For a license part number ending in "K9", the TLS proxy limit depends on your configuration and the platform model.
	Note K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted.
	You might also use SRTP encryption sessions for your connections:
	• For a K8 license, SRTP sessions are limited to 250.
	• For a K9 license, there is no limit.
	Note Only calls that require encryption/decryption for media are counted toward the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count toward the limit.

Guidelines and Limitations

Context Mode Guidelines

Supported in single context mode only.

Firewall Mode Guidelines

Supported in routed firewall mode only.

IPv6 Guidelines

Does not support IPv6 addresses.

Additional Guidelines and Limitations

Cisco Intercompany Media Engine has the following limitations:

- Fax is not supported. Fax capability needs to be disabled on the SIP trunk.
- Stateful failover of Cisco Unified Intercompany Media Engine is not supported. During failover, existing calls traversing the Cisco Intercompany Media Engine Proxy disconnect; however, new calls successfully traverse the proxy after the failover completes.

- Having Cisco UCMs on more than one of the ASA interfaces is not supported with the Cisco Intercompany Media Engine Proxy. Having the Cisco UCMs on one trusted interface is especially necessary in an off path deployment because the ASA requires that you specify the listening interface for the mapping service and the Cisco UCMs must be connected on one trusted interface.

- Multipart MIME is not supported.

- Only existing SIP features and messages are supported.

- H.264 is not supported.

- RTCP is not supported. The ASA drops any RTCP traffic sent from the inside interface to the outside interface. The ASA does not convert RTCP traffic from the inside interface into SRTP traffic.

- The Cisco Intercompany Media Engine Proxy configured on the ASA creates a dynamic SIP trunk for each connection to a remote enterprise. However, you cannot configure a unique subject name for each SIP trunk. The Cisco Intercompany Media Engine Proxy can have only one subject name configured for the proxy.

  Additionally, the subject DN you configure for the Cisco Intercompany Media Engine Proxy match the domain name that has been set for the local Cisco UCM.

- If a service policy rule for the Cisco Intercompany Media Engine Proxy is removed (by using the no service policy command) and reconfigured, the first call traversing the ASA will fail. The call fails over to the PSTN because the Cisco UCM does not know the connections are cleared and tries to use the recently cleared IME SIP trunk for the signaling.

  To resolve this issue, you must additionally enter the **clear connection all** command and restart the ASA. If the failure is due to failover, the connections from the primary ASA are not synchronized to the standby ASA.

- After the **clear connection all** command is issued on an ASA enabled with a UC-IME Proxy and the IME call fails over to the PSTN, the next IME call between an originating and terminating SCCP IP phone completes but does not have audio and is dropped after the signaling session is established.

  An IME call between SCCP IP phones use the IME SIP trunk in both directions. Namely, the signaling from the calling to called party uses the IME SIP trunk. Then, the called party uses the reverse IME SIP trunk for the return signaling and media exchange. However, this connection is already cleared on the ASA, which causes the IME call to fail.

  The next IME call (the third call after the **clear connection all** command is issued), will be completely successful.

  Note   This limitation does not apply when the originating and terminating IP phones are configured with SIP.

- The ASA must be licensed and configured with enough TLS proxy sessions to handle the IME call volume. See Licensing for Cisco Intercompany Media Engine, page 19-8 for information about the licensing requirements for TLS proxy sessions.

  This limitation occurs because an IME call cannot fall back to the PSTN when there are not enough TLS proxy sessions left to complete the IME call. An IME call between two SCCP IP phones requires the ASA to use two TLS proxy sessions to successfully complete the TLS handshake.

  Assume for example, the ASA is configured to have a maximum of 100 TLS proxy sessions and IME calls between SCCP IP phones establish 101 TLS proxy sessions. In this example, the next IME call is initiated successfully by the originating SCCP IP phone but fails after the call is accepted by the terminating SCCP IP phone. The terminating IP phone rings and on answering the call, the call hangs due to an incomplete TLS handshake. The call does not fall back to the PSTN.

Configuring Cisco Intercompany Media Engine Proxy

This section contains the following topics:

Task Flow for Configuring Cisco Intercompany Media Engine

Figure 19-6 provides an example for a basic deployment of the Cisco Intercompany Media Engine. The following tasks include command line examples based on Figure 19-6.

Figure 19-6   Example for Basic (in-line) Deployment Tasks
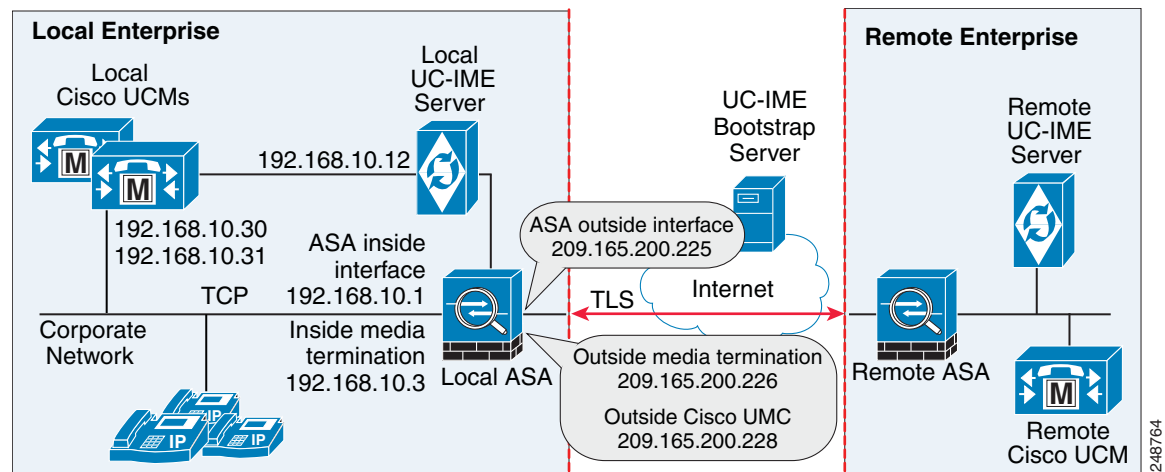
Note   Step 1 through Step 8 apply to both basic (in-line) and off path deployments and Step 9 applies only to off path deployment.

To configure a Cisco Intercompany Media Engine for a basic deployment, perform the following tasks.

Step 1   Configure static NAT for Cisco UCM. See Configuring NAT for Cisco Intercompany Media Engine Proxy, page 19-12.

Or

Configure PAT for the UCM server. See Configuring PAT for the Cisco UCM Server, page 19-14.

Step 2   Create ACLs for Cisco Intercompany Media Engine Proxy. See Creating ACLs for Cisco Intercompany Media Engine Proxy, page 19-16.

Step 3   Create the media termination address instance for Cisco Intercompany Media Engine Proxy. See Creating the Media Termination Instance, page 19-17.

Step 4   Create the Cisco Intercompany Media Engine Proxy. See Creating the Cisco Intercompany Media Engine Proxy, page 19-19.

Step 5   Create trustpoints and generate certificates for the Cisco Intercompany Media Engine Proxy. See Creating Trustpoints and Generating Certificates, page 19-22.

Step 6   Create the TLS proxy. See Creating the TLS Proxy, page 19-25.

Step 7   Configure SIP inspection for the Cisco Intercompany Media Engine Proxy. See Enabling SIP Inspection for the Cisco Intercompany Media Engine Proxy, page 19-26.

Step 8   (Optional) Configure TLS within the enterprise. See (Optional) Configuring TLS within the Local Enterprise, page 19-28.

Step 9   (Optional) Configure off path signaling. See (Optional) Configuring Off Path Signaling, page 19-31.

> **Note**   You only perform Step 9 when you are configuring the Cisco Intercompany Media Engine Proxy in an off path deployment.

Configuring NAT for Cisco Intercompany Media Engine Proxy

To configure auto NAT, you first configure an object; then use the **nat** command in the object configuration mode.

The example command lines in this task are based on a basic (in-line) deployment. See Figure 19-6 on page 19-11 for an illustration explaining the example command lines in this task.

Alternatively, you can configure PAT for the Cisco Intercompany Media Engine Proxy. See Configuring PAT for the Cisco UCM Server, page 19-14.

Figure 19-7   Example for Configuring NAT for a Deployment
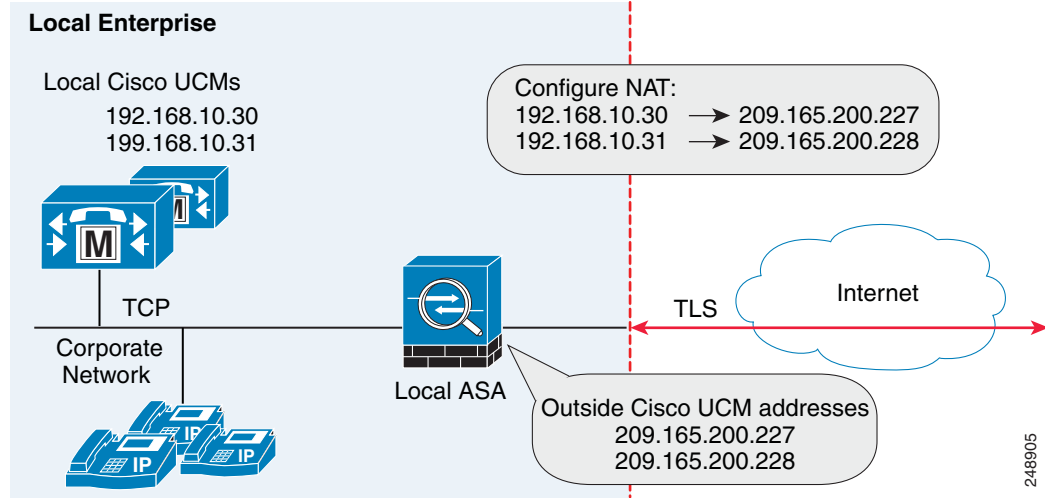
To configure auto NAT rules for the Cisco UCM server, perform the following steps:

	Command	Purpose
Step 1	hostname(config)# **object network** *name* **Example:** hostname(config)# object network ucm_real_192.168.10.30 hostname(config)# object network ucm_real_192.168.10.31	Configures a network object for the real address of Cisco UCM that you want to translate.
Step 2	hostname(config-network-object)# **host** *ip_address* **Example:** hostname(config-network-object)# host 192.168.10.30 hostname(config-network-object)# host 192.168.10.31	Specifies the real IP address of the Cisco UCM host for the network object.
Step 3	(Optional) hostname(config-network-object)# **description** *string* **Example:** hostname(config-network-object)# description "Cisco UCM Real Address"	Provides a description of the network object.
Step 4	hostname(config-network-object)# **exit**	Exits from the objects configuration mode.
Step 5	hostname(config)# **object network** *name* **Example:** hostname(config)# object network ucm_map_209.165.200.228	Configures a network object for the mapped address of the Cisco UCM.

	Command	Purpose

Step 6	`hostname(config-network-object)# host ip_address`
	
	Example:
	`hostname(config-network-object)# host 209.165.200.228`
		Specifies the mapped IP address of the Cisco UCM host for the network object.

Step 7	(Optional)
	
	`hostname(config-network-object)# description string`
	
	Example:
	`hostname(config-network-object)# description "Cisco UCM Mapped Address"`
		Provides a description of the network object.

Step 8	`hostname(config-network-object)# exit`	Exits from the objects configuration mode.

Step 9	`hostname(config)# nat (inside,outside) source static real_obj mapped_obj`
	
	Example:
	`hostname(config)# nat (inside,outside) source static ucm_real_192.168.10.30 ucm_209.165.200.228`
	`hostname(config)# nat (inside,outside) source static ucm_real_192.168.10.31 ucm_209.165.200.228`
		Specifies the address translation on the network objects created in this procedure.
		
		Where *real_obj* is the *name* that you created in Step 1 in this task.
		
		Where *mapped_obj* is the *name* that you created in Step 5 in this task.

What to Do Next

Create the ACLs for the Cisco Intercompany Media Engine Proxy. See Creating ACLs for Cisco Intercompany Media Engine Proxy, page 19-16.

Configuring PAT for the Cisco UCM Server

Perform this task as an alternative to configuring NAT for the Cisco Intercompany Media Engine Proxy.

Figure 19-8   Example for Configuring PAT for a Deployment

> **Note**   You only perform this step when NAT is not configured for the Cisco UCM server.

To configure PAT for the Cisco UCM server, perform the following steps:

	Command	Purpose

Step 1	`hostname(config)# `**`object network`**` name`
	
	Example:
	`hostname(config)# object network`
	`ucm-pat-209.165.200.228`
		Configures a network object for the outside IP address of Cisco UCM that you want to translate.

Step 2	`hostname(config-network-object)# `**`host`**` ip_address`
	
	Example:
	`hostname(config-network-object)# host`
	`209.165.200.228`
		Specifies the real IP address of the Cisco UCM host for the network object.

Step 3	`hostname(config-network-object)# `**`exit`**	Exits from the objects configuration mode.

Step 4	`hostname(config)# `**`object service`**` name`
	
	Example:
	`hostname(config)# object service tcp_5070`
	`hostname(config)# object service tcp_5071`
		Creates a service object for the outside Cisco Intercompany Media Engine port.

Step 5	`hostname(config-service-object)# `**`tcp source eq`**` port`
	
	Example:
	`hostname(config-service-object)# tcp source eq 5070`
	`hostname(config-service-object)# tcp source eq 5071`
		Specifies the port number.

Step 6	`hostname(config-service-object)# `**`exit`**	Exits from the objects configuration mode.

Step 7	`hostname(config)# `**`object network`**` name`
	
	Example:
	`hostname(config)# object network`
	`ucm-real-192.168.10.30`
	`hostname(config)# object network`
	`ucm-real-192.168.10.31`
		Configures a network object to represent the real IP address of Cisco UCM.

Step 8	`hostname(config-network-object)# `**`host`**` ip_address`
	
	Example:
	`hostname(config-network-object)# host 192.168.10.30`
	`hostname(config-network-object)# host 192.168.10.31`
		Specifies the real IP address of the Cisco UCM host for the network object.

Step 9	`hostname(config-network-object)# `**`exit`**	Exits from the objects configuration mode.

Step 10	`hostname(config)# `**`object service`**` name`
	
	Example:
	`hostname(config)# object service tcp_5570`
	`hostname(config)# object service tcp_5571`
		Creates a service objects for Cisco UCM SIP port.

	Command	Purpose

Step 11	`hostname(config-service-object)# ` **`tcp source eq`** *`port`*
	
	Example:
	`hostname(config-service-object)# tcp source eq 5570`
	`hostname(config-service-object)# tcp source eq 5571`
		Specifies the port number.

Step 12	`hostname(config-service-object)# ` **`exit`**	Exits from the objects configuration mode.

Step 13	`hostname(config)# ` **`nat (inside,outside) source static`** *`real_obj mapped_obj`* **`service`** *`real_port mapped_port`*
	
	Example:
	`hostname(config)# nat (inside,outside) source static ucm-real-192.168.10.30 ucm-pat-209.165.200.228 service tcp_5070 tcp_5570`
	`hostname(config)# nat (inside,outside) source static ucm-real-192.168.10.31 ucm-pat-128.106.254.5 service tcp_5071 tcp_5571`
		Creates a static mapping for Cisco UCM.
		
		Where real_obj is the name that you created in Step 1 in this task.
		
		Where mapped_obj is the name that you created in Step 7 in this task.
		
		Where real_port is the name that you created in Step 4 in this task.
		
		Where mapped_obj is the name that you created in Step 10 in this task.

Creating ACLs for Cisco Intercompany Media Engine Proxy

To configure ACLs for the Cisco Intercompany Media Engine Proxy to reach the Cisco UCM server, perform the following steps.

The example command lines in this task are based on a basic (in-line) deployment. See Figure 19-6 on page 19-11 for an illustration explaining the example command lines in this task.

	Command	Purpose
Step 1	`hostname(config)# ` **`access-list`** *`id`* **`extended permit tcp any host`** *`ip_address`* **`eq`** *`port`* **Example:** `hostname(config)# access-list incoming extended permit tcp any host 192.168.10.30 eq 5070`	Adds an Access Control Entry (ACE). An ACL is made up of one or more ACEs with the same ACL ID. This ACE provides access control by allowing incoming access for Cisco Intercompany Media Engine connections on the specified port. In the *ip_address* argument, provide the real IP address of Cisco UCM.
Step 2	`hostname(config)# ` **`access-group`** *`access-list`* **`in interface`** *`interface_name`* **Example:** `hostname(config)# access-group incoming in interface outside`	Binds the ACL to an interface.

	Command	Purpose

Step 3	`hostname(config)# access-list id extended permit tcp any host ip_address eq port`
	
	Example:
	`hostname(config)# access-list ime-inbound-sip extended permit tcp any host 192.168.10.30 eq 5070`
		Adds an ACE. This ACE allows the ASA to allow inbound SIP traffic for Cisco Intercompany Media Engine. This entry is used to classify traffic for the class and policy map.
		
		Note The port that you configure here must match the trunk settings configured on Cisco UCM. See the Cisco Unified Communications Manager documentation for information about this configuration setting.

Step 4	`hostname(config)# access-list id extended permit tcp ip_address mask any range range`
	
	Example:
	`hostname(config)# access-list ime-outbound-sip extended permit tcp 192.168.10.30 255.255.255.255 any range 5000 6000`
		Adds an ACE. This ACE allows the ASA to allow outbound SIP traffic for Cisco Intercompany Media Engine (in the example, any TCP traffic with source as 192.168.10.30 and destination port range between 5000 and 6000). This entry is used to classify traffic for the class and policy map.
		
		Note Ensure that TCP traffic between Cisco UCM and the Cisco Intercompany Media Engine server does not use this port range (if that connection goes through the ASA).

Step 5	`hostname(config)# access-list id permit tcp any host ip_address eq 6084`
	
	Example:
	`hostname(config)# access-list ime-traffic permit tcp any host 192.168.10.12 eq 6084`
		Adds an ACE. This ACE allows the ASA to allow traffic from the Cisco Intercompany Media Engine server to remote Cisco Intercompany Media Engine servers.

Step 6	`hostname(config)# access-list id permit tcp any host ip_address eq 8470`
	
	Example:
	`hostname(config)# access-list ime-bootserver-traffic permit tcp any host 192.168.10.12 eq 8470`
		Adds an ACE. This ACE allows the ASA to allow traffic from the Cisco Intercompany Media Engine server to the Bootstrap server for the Cisco Intercompany Media Engine.

What to Do Next

Create the media termination instance on the ASA for the Cisco Intercompany Media Engine Proxy. See Creating the Media Termination Instance, page 19-17.

Creating the Media Termination Instance

Guidelines

The media termination address you configure must meet these requirements:

- If you decide to configure a media-termination address on interfaces (rather than using a global interface), you must configure a media-termination address on at least two interfaces (the inside and an outside interface) before applying the service policy for the Cisco Intercompany Media Engine Proxy. Otherwise, you will receive an error message when enabling the proxy with SIP inspection.

> **Note**   Cisco recommends that you configure the media-termination address for the Cisco Intercompany Media Engine Proxy on interfaces rather than configuring a global media-termination address.

- The Cisco Intercompany Media Engine Proxy can use only one type of media termination instance at a time; for example, you can configure a global media-termination address for all interfaces or configure a media-termination address for different interfaces. However, you cannot use a global media-termination address and media-termination addresses configured for each interface at the same time.

> **Note**   If you change any Cisco Intercompany Media Engine Proxy settings after you create the media-termination address for the proxy, you must reconfigure the media-termination address by using the **no media-termination** command, and then reconfiguring it as described in this procedure.

Procedure

Create the media termination instance to use with the Cisco Intercompany Media Engine Proxy.

The example command lines in this task are based on a basic (in-line) deployment. See Figure 19-6 on page 19-11 for an illustration explaining the example command lines in this task.

To create the media termination instance for the Cisco Intercompany Media Engine Proxy, perform the following steps:

	Command	Purpose
Step 1	hostname(config)# **media-termination** *instance_name* **Example:** hostname(config)# **media-termination** *uc-ime-media-term*	Creates the media termination instance that you attach to the Cisco Intercompany Media Engine Proxy.
Step 2	hostname(config-media-termination)# **address** *ip_address* **interface** *intf_name* **Examples:** hostname(config-media-termination)# address 209.165.200.228 interface outside	Configures the media-termination address used by the outside interface of the ASA. The outside IP address must be a publicly routable address that is an unused IP address within the address range on that interface. See Creating the Cisco Intercompany Media Engine Proxy, page 19-19 for information about the UC-IME proxy settings. See CLI configuration guide for information about the **no service-policy** command.

	Command	Purpose
Step 3	hostname(config-media-termination)# **address** *ip_address* **interface** *intf_name* **Examples:** hostname(config-media-termination)# address 192.168.10.3 interface inside	Configures a media termination address used by the inside interface of the ASA. **Note** The IP address must be an unused IP address within the same subnet on that interface.
Step 4	(Optional) hostname(config-media-termination)# ***rtp-min-port*** *port1* ***rtp-maxport*** *port2* **Examples:** hostname(config-media-termination)# rtp-min-port 1000 rtp-maxport 2000	Configures the rtp-min-port and rtp-max-port limits for the Cisco Intercompany Media Engine Proxy. Configure the RTP port range for the media termination point when you need to scale the number of calls that the Cisco Intercompany Media Engine supports. Where *port1* specifies the minimum value for the RTP port range for the media termination point, where port1 can be a value from 1024 to 65535. By default, the value for *port1* is 16384. Where *port2* specifies the maximum value for the RTP port range for the media termination point, where port2 can be a value from 1024 to 65535. By default, the value for *port2* is 32767.

What To Do Next

Once you have created the media termination instance, create the Cisco Intercompany Media Engine Proxy. See Creating the Cisco Intercompany Media Engine Proxy, page 19-19.

Creating the Cisco Intercompany Media Engine Proxy

To create the Cisco Intercompany Media Engine Proxy, perform the following steps.

The example command lines in this task are based on a basic (in-line) deployment. See Figure 19-6 on page 19-11 for an illustration explaining the example command lines in this task.

Note   You cannot change any of the configuration settings for the Cisco Intercompany Media Engine Proxy described in this procedure when the proxy is enabled for SIP inspection. Remove the Cisco Intercompany Media Engine Proxy from SIP inspection before changing any of the settings described in this procedure.

	Command	Purpose
Step 1	hostname(config)# **uc-ime** *uc_ime_name* **Example:** hostname(config)# uc-ime local-ent-ime	Configures the Cisco Intercompany Media Engine Proxy. Where *uc_ime_name* is the name of the Cisco Intercompany Media Engine Proxy. The name is limited to 64 characters. Only one Cisco Intercompany Media Engine Proxy can be configured on the ASA.
Step 2	hostname(config-uc-ime)# **media-termination** *mta_instance_name* **Example:** hostname(config-uc-ime)# media-termination ime-media-term	Specifies the media termination instance used by the Cisco Intercompany Media Engine Proxy. **Note** You must create the media termination instance before you specify it in the Cisco Intercompany Media Engine Proxy. Where *mta_instance_name* is the *instance_name* that you created in Step 1 of Creating the Media Termination Instance. See Creating the Media Termination Instance, page 19-17 for the steps to create the media termination instance.
Step 3	hostname(config-uc-ime)# **ucm address** *ip_address* **trunk-security-mode** [**nonsecure** | **secure**] **Example:** hostname(config-uc-ime)# ucm address 192.168.10.30 trunk-security-mode non-secure	Specifies the Cisco UCM server in the enterprise. You must specify the real IP address of the Cisco UCM server. Do not specify a mapped IP address for the server. **Note** You must include an entry for each Cisco UCM in the cluster with Cisco Intercompany Media Engine that has a SIP trunk enabled. Where the **nonsecure** and **secure** options specify the security mode of the Cisco UCM or cluster of Cisco UCMs. **Note** Specifying **secure** for Cisco UCM or Cisco UCM cluster indicates that Cisco UCM or Cisco UCM cluster is initiating TLS; therefore, you must configure TLS for components. See (Optional) Configuring TLS within the Local Enterprise, page 19-28. You can specify the **secure** option in this task or you can update it later while configuring TLS for the enterprise. See Step 11 in (Optional) Configuring TLS within the Local Enterprise, page 19-28.

	Command	Purpose
Step 4	`hostname(config-uc-ime)# ` **`ticket epoch`** `n` **`password`** `password`	Configures the ticket epoch and password for Cisco Intercompany Media Engine.
		Where *n* is an integer from 1-255. The epoch contains an integer that updates each time that the password is changed. When the proxy is configured the first time and a password entered for the first time, enter 1 for the epoch integer. Each time you change the password, increment the epoch to indicate the new password. You must increment the epoch value each time your change the password.
	Example: `hostname(config-uc-ime)# ticket epoch 1 password password1234`	Typically, you increment the epoch sequentially; however, the ASA allows you to choose any value when you update the epoch.
		If you change the epoch value, the current password is invalidated and you must enter a new password.
		Where *password* contains a minimum of 10 and a maximum of 64 printable character from the US-ASCII character set. The allowed characters include 0x21 to 0x73 inclusive, and exclude the space character.
		We recommend a password of at least 20 characters. Only one password can be configured at a time.
		The ticket password is stored onto flash. The output of the **show running-config uc-ime** command displays ***** instead of the password string.
		Note The epoch and password that you configure on the ASA must match the epoch and password configured on the Cisco Intercompany Media Engine server. See the Cisco Intercompany Media Engine server documentation for information.

	Command	Purpose
Step 5	(Optional) `hostname(config-uc-ime)# fallback monitoring timer` `timer_millisec	hold-down timer timer_sec`  **Example:** `hostname(config-uc-ime)# fallback monitoring timer` `120` `hostname(config-uc-ime)# fallback hold-down timer 30`	Specifies the fallback timers for Cisco Intercompany Media Engine. Specifying **monitoring timer** sets the time between which the ASA samples the RTP packets received from the Internet. The ASA uses the data sample to determine if fallback to the PSTN is needed for a call. Where *timer_millisec* specifies the length of the monitoring timer. By default, the length is 100 milliseconds for the monitoring timer and the allowed range is 10-600 ms. Specifying **hold-down timer** sets the amount of time that ASA waits before notifying Cisco UCM whether to fall back to PSTN. Where *timer_sec* specifies the length of the hold-down timer. By default, the length is 20 seconds for the hold-down timer and the allowed range is 10-360 seconds. If you do not use this command to specify fallback timers, the ASA uses the default settings for the fallback timers.

Step 6	(Optional)
	
	`hostname(config-uc-ime)# fallback sensitivity-file`
	`file_name`
	
	Example:
	`hostname(config-uc-ime)# fallback sensitivity-file`
	`ime-fallback-sensitvity.fbs`
		Specifies the file to use for mid-call PSTN fallback.
		
		Where *file_name* must be the name of a file on disk that includes the .fbs file extension.
		
		The fallback file is used to determine whether the QoS of the call is poor enough for the Cisco Intercompany Media Engine to move the call to the PSTN.

What to Do Next

Install the certificate on the local entity truststore. You could also enroll the certificate with a local CA trusted by the local entity.

Creating Trustpoints and Generating Certificates

You need to generate the keypair for the certificate used by the ASA, and configure a trustpoint to identify the certificate sent by the ASA in the TLS handshake.

The example command lines in this task are based on a basic (in-line) deployment. See Figure 19-6 on page 19-11 for an illustration explaining the example command lines in this task.

Note   This task instructs you on how to create trustpoints for the local enterprise and the remote enterprise and how to exchange certificates between these two enterprises. This task does not provide steps for creating trustpoints and exchanging certificates between the local Cisco UCM and the local ASA. However, if you require additional security within the local enterprise, you must perform the optional task (Optional)

Configuring TLS within the Local Enterprise, page 19-28. Performing that task allows for secure TLS connections between the local Cisco UCM and the local ASA. The instructions in that task describe how to create trustpoints between the local Cisco UCM and the local ASA.

Prerequisites for Installing Certificates

To create a proxy certificate on the ASA that is trusted by the remote entity, obtain a certificate from a trusted CA or export it from the remote enterprise ASA.

To export the certificate from the remote enterprise, you enter the following command on the remote ASA:

hostname(config)# **crypto ca export trustpoint identity-certificate**

The ASA prompts displays the certificate in the terminal screen. Copy the certificate from the terminal screen. You will need the certificate text in Step 5 of this task.

Procedure

To create the trustpoints and generate certificates, perform the following steps:

	Command	**Purpose**
Step 1	hostname(config)# **crypto key generate rsa label** *key-pair-label* **modulus** *size* **Example:** hostname(config)# crypto key generate rsa label local-ent-key modulus 2048	On the local ASA, creates the RSA keypair that can be used for the trustpoints. This is the keypair and trustpoint for the local entities signed certificate. The modulus key size that you select depends on the level of security that you want to configure and on any limitations imposed by the CA from which you are obtaining the certificate. The larger the number that you select, the higher the security level will be for the certificate. Most CAs recommend 2048 for the key modulus size; however, **Note** GoDaddy requires a key modulus size of 2048.
Step 2	hostname(config)# **crypto ca trustpoint** *trustpoint_name* **Example:** hostname(config)# crypto ca trustpoint local_ent	Enters the trustpoint configuration mode for the specified trustpoint so that you can create the trustpoint for the local entity. A trustpoint represents a CA identity and possibly a device identity, based on a certificate issued by the CA. Maximum name length is 128 characters.
Step 3	hostname(config-ca-trustpoint)# **subject-name** *X.500_name* **Example:** hostname(config-ca-trustpoint)# subject-name cn=Ent-local-domain-name**	Includes the indicated subject DN in the certificate during enrollment. **Note** The domain name that you enter here must match the domain name that has been set for the local Cisco UCM. For information about how to configure the domain name for Cisco UCM, see the Cisco Unified Communications Manager documentation for information.

	Command	Purpose
Step 4	hostname(config-ca-trustpoint)# **keypair** *keyname* **Example:** hostname(config-ca-trustpoint)# keypair local-ent-key	Specifies the key pair whose public key is to be certified.
Step 5	hostname(config-ca-trustpoint)# **enroll terminal**	Specifies that you will use the "copy and paste" method of enrollment with this trustpoint (also known as manual enrollment).
Step 6	hostname(config-ca-trustpoint)# **exit**	Exits from the CA Trustpoint configuration mode.
Step 7	hostname(config)# **crypto ca enroll** *trustpoint* **Example:** hostname(config)# crypto ca enroll remote-ent % % Start certificate enrollment ... % The subject name in the certificate will be: % cn=enterpriseA % The fully-qualified domain name in the certificate will @ be: ciscoasa % Include the device serial number in the subject name? [yes/no]: no Display Certificate Request to terminal? [yes/no]: yes	Starts the enrollment process with the CA. Where *trustpoint* is the same as the value you entered for *trustpoint_name* in Step 2. When the trustpoint is configured for manual enrollment (**enroll terminal** command), the ASA writes a base-64-encoded PKCS10 certification request to the console and then displays the CLI prompt. Copy the text from the prompt. Submit the certificate request to the CA, for example, by pasting the text displayed at the prompt into the certificate signing request enrollment page on the CA website. When the CA returns the signed identity certificate, proceed to Step 8 in this procedure.
Step 8	hostname(config)# **crypto ca import** *trustpoint* **certificate** **Example:** hostname(config)# crypto ca import remote-ent certificate	Imports the signed certificate received from the CA in response to a manual enrollment request. Where *trustpoint* specifies the trustpoint you created in Step 2. The ASA prompts you to paste the base-64 formatted signed certificate onto the terminal.
Step 9	hostname(config)# **crypto ca authenticate** *trustpoint* **Example:** hostname(config)# crypto ca authenticate remote-ent	Authenticates the third-party identity certificate received from the CA. The identity certificate is associated with a trustpoint created for the remote enterprise. The ASA prompts you to paste the base-64 formatted identity certificate from the CA onto the terminal.

What to Do Next

Create the TLS proxy for the Cisco Intercompany Media Engine. See Creating the TLS Proxy, page 19-25.

Creating the TLS Proxy

Because either enterprise, namely the local or remote Cisco UCM servers, can initiate the TLS handshake (unlike IP Telephony or Cisco Mobility Advantage, where only the clients initiate the TLS handshake), you must configure by-directional TLS proxy rules. Each enterprise can have an ASA as the TLS proxy.

Create TLS proxy instances for the local and remote entity initiated connections respectively. The entity that initiates the TLS connection is in the role of "TLS client." Because the TLS proxy has a strict definition of "client" and "server" proxy, two TLS proxy instances must be defined if either of the entities could initiate the connection.

The example command lines in this task are based on a basic (in-line) deployment. See Figure 19-6 on page 19-11 for an illustration explaining the example command lines in this task.

To create the TLS proxy, perform the following steps:

	Command	**Purpose**
Step 1	hostname(config)# **tls-proxy** proxy_name **Example:** hostname(config)# tls-proxy local_to_remote-ent	Creates the TLS proxy for the outbound connections.
Step 2	hostname(config-tlsp)# **client trust-point** *proxy_trustpoint* **Example:** hostname(config-tlsp)# client trust-point local-ent	For **outbound** connections, specifies the trustpoint and associated certificate that the adaptive security appliance uses in the TLS handshake when the adaptive security appliance assumes the role of the TLS client. The certificate must be owned by the adaptive security appliance (identity certificate). Where *proxy_trustpoint* specifies the trustpoint defined by the **crypto ca trustpoint** command in Step 2 in Creating Trustpoints and Generating Certificates, page 19-22.
Step 3	hostname(config-tlsp)# **client cipher-suite** *cipher_suite* **Example:** hostname(config-tlsp)# client cipher-suite aes128-sha1 aes256-sha1 3des-sha1 null-sha1	For outbound connections, controls the TLS handshake parameter for the cipher suite. Where *cipher_suite* includes des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, or null-sha1. For client proxy (the proxy acts as a TLS client to the server), the user-defined cipher suite replaces the default cipher suite, or the one defined by the **ssl encryption** command. Use this command to achieve difference ciphers between the two TLS sessions. You should use AES ciphers with the Cisco UCM server.
Step 4	hostname(config-tlsp)# **exit**	Exits from the TLS proxy configuration mode.
Step 5	hostname(config)# **tls-proxy** *proxy_name* **Example:** hostname(config)# tls-proxy remote_to_local-ent	Create the TLS proxy for inbound connections.

	Command	Purpose
Step 6	hostname(config-tlsp)# **server trust-point** *proxy_trustpoint* **Example:** hostname(config-tlsp)# server trust-point local-ent	For **inbound** connections, specifies the proxy trustpoint certificate presented during TLS handshake. The certificate must be owned by the adaptive security appliance (identity certificate). Where *proxy_trustpoint* specifies the trustpoint defined by the **crypto ca trustpoint** command in Step 2 in Creating Trustpoints and Generating Certificates, page 19-22. Because the TLS proxy has strict definition of client proxy and server proxy, two TLS proxy instances must be defined if either of the entities could initiate the connection.
Step 7	hostname(config-tlsp)# **client cipher-suite** *cipher_suite* **Example:** hostname(config-tlsp)# client cipher-suite aes128-sha1 aes256-sha1 3des-sha1 null-sha1	For inbound connections, controls the TLS handshake parameter for the cipher suite. Where *cipher_suite* includes des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, or null-sha1.
Step 8	hostname(config-tlsp)# **exit**	Exits from the TSL proxy configuration mode.
Step 9	hostname(config)# **ssl encryption 3des-sha1 aes128-sha1** [*algorithms*]	Specifies the encryption algorithms that the SSL/TLS protocol uses. Specifying the 3des-sha1 and aes128-sha1 is required. Specifying other algorithms is optional. **Note** The Cisco Intercompany Media Engine Proxy requires that you use strong encryption. You must specify this command when the proxy is licensed using a K9 license.

What to Do Next

Once you have created the TLS proxy, enable it for SIP inspection.

Enabling SIP Inspection for the Cisco Intercompany Media Engine Proxy

Enable the TLS proxy for SIP inspection and define policies for both entities that could initiate the connection.

The example command lines in this task are based on a basic (in-line) deployment. See Figure 19-6 on page 19-11 for an illustration explaining the example command lines in this task.

Note   If you want to change any Cisco Intercompany Media Engine Proxy settings after you enable SIP inspection, you must enter the **no service-policy** command, and then reconfigure the service policy as described in this procedure. Removing and reconfiguring the service policy does not affect existing calls; however, the first call traversing the Cisco Intercompany Media Engine Proxy will fail. Enter the **clear connection** command and restart the ASA.

To enable SIP inspection for the Cisco Intercompany Media Engine Proxy, perform the following steps:

	Command	Purpose
Step 1	hostname(config)# **class-map** *class_map_name* **Example:** hostname(config)# class-map ime-inbound-sip	Defines a class for the inbound Cisco Intercompany Media Engine SIP traffic.
Step 2	hostname(config-cmap)# **match access-list** *access_list_name* **Example:** hostname(config-cmap)# match access-list ime-inbound-sip	Identifies the SIP traffic to inspect. Where the *access_list_name* is the ACL you created in Step 3, page 19-17 of the task Creating ACLs for Cisco Intercompany Media Engine Proxy.
Step 3	hostname(config-cmap)# **exit**	Exits from the class map configuration mode.
Step 4	hostname(config)# **class-map** *class_map_name* **Example:** hostname(config)# class-map ime-outbound-sip	Defines a class for the outbound SIP traffic from Cisco Intercompany Media Engine.
Step 5	hostname(config)# **match access-list** *access_list_name* **Example:** hostname(config-cmap)# match access-list ime-outbound-sip	Identifies which outbound SIP traffic to inspect. Where the *access_list_name* is the ACL you created in Step 4, page 19-17 of the task Creating ACLs for Cisco Intercompany Media Engine Proxy.
Step 6	hostname(config-cmap)# **exit**	Exits from the class map configuration mode.
Step 7	hostname(config)# **policy-map** *name* **Example:** hostname(config)# policy-map ime-policy	Defines the policy map to which to attach the actions for the class of traffic.
Step 8	hostname(config-pmap)# **class** *classmap_name* **Example:** hostname(config-pmap)# class ime-outbound-sip	Assigns a class map to the policy map so that you can assign actions to the class map traffic. Where *classmap_name* is the name of the SIP class map that you created in Step 1 in this task.
Step 9	hostname(config-pmap-c)# **inspect sip** [*sip_map*] **tls-proxy** *proxy_name* **uc-ime** *uc_ime_map* **Example:** hostname(config-pmap-c)# inspect sip tls-proxy local_to_remote-ent uc-ime local-ent-ime	Enables the TLS proxy and Cisco Intercompany Media Engine Proxy for the specified SIP inspection session.
Step 10	hostname(config-cmap-c)# **exit**	Exits from the policy map class configuration mode.
Step 11	hostname(config-pmap)# **class** *class_map_name* **Example:** hostname(config-pmap)# class ime-inbound-sip	Assigns a class map to the policy map so that you can assign actions to the class map traffic. Where *classmap_name* is the name of the SIP class map that you created in Step 4 in this task.

	Command	Purpose
Step 12	hostname(config-pmap-c)# **inspect sip** [*sip_map*] **tls-proxy** *proxy_name* **uc-ime** *uc_ime_map* **Example:** hostname(config-pmap-c)# inspect sip tls-proxy remote-to-local-ent uc-ime local-ent-ime	Enables the TLS proxy and Cisco Intercompany Media Engine Proxy for the specified SIP inspection session.
Step 13	hostname(config-pmap-c)# **exit**	Exits from the policy map class configuration mode.
Step 14	hostname(config-pmap)# **exit**	Exits from the policy map configuration mode.
Step 15	hostname(config)# **service-policy** *policymap_name* **global** **Example:** hostname(config)# service-policy ime-policy global	Enables the service policy for SIP inspection for all interfaces. Where *policymap_name* is the name of the policy map you created in Step 7 of this task. See Creating the Cisco Intercompany Media Engine Proxy, page 19-19 for information about the UC-IME proxy settings. See CLI configuration guide for information about the **no service-policy** command.

What to Do Next

Once you have enabled the TLS proxy for SIP inspection, if necessary, configure TLS within the enterprise. See (Optional) Configuring TLS within the Local Enterprise, page 19-28.

(Optional) Configuring TLS within the Local Enterprise

This task is not required if TCP is allowable within the inside network.

TLS within the enterprise refers to the security status of the Cisco Intercompany Media Engine trunk as seen by the ASA.

Note If the transport security for the Cisco Intercompany Media Engine trunk changes on Cisco UCM, it must be changed on the ASA as well. A mismatch will result in call failure. The ASA does not support SRTP with non-secure IME trunks. The ASA assumes SRTP is allowed with secure trunks. So 'SRTP Allowed' must be checked for IME trunks if TLS is used. The ASA supports SRTP fallback to RTP for secure IME trunk calls.

Prerequisites

On the local Cisco UCM, download the Cisco UCM certificate. See the Cisco Unified Communications Manager documentation for information. You will need this certificate when performing Step 6 of this procedure.

Procedure

To configure TLS within the local enterprise, perform the following steps on the local ASA:

	Commands	Purpose

Step 1	`hostname(config)# `**`crypto key generate rsa label`** `key-pair-label`
	`hostname(config)# `**`crypto ca trustpoint`** `trustpoint_name`
	`hostname(config-ca-trustpoint)# `**`enroll self`**
	`hostname(config-ca-trustpoint)# `**`keypair`** `keyname`
	`hostname(config-ca-trustpoint)# `**`subject-name`** `x.500_name`
	`Example:`
	`hostname(config)# crypto key generate rsa label local-ent-key`
	`hostname(config)# crypto ca trustpoint local-asa`
	`hostname(config-ca-trustpoint)# enroll self`
	`hostname(config-ca-trustpoint)# keypair key-local-asa`
	`hostname(config-ca-trustpoint)# subject-name cn=Ent-local-domain-name**., o="Example Corp"`
		Creates an RSA key and trustpoint for the self-signed certificate.
		
		Where `key-pair-label` is the RSA key for the local ASA.
		
		Where `trustpoint_name` is the trustpoint for the local ASA.
		
		Where *keyname* is key pair for the local ASA.
		
		Where `x.500_name` includes the X.500 distinguished name of the local ASA; for example, `cn=Ent-local-domain-name**.`
		
		Note The domain name that you enter here must match the domain name that has been set for the local Cisco UCM. For information about how to configure the domain name for Cisco UCM, see the Cisco Unified Communications Manager documentation for information.

Step 2	`hostname(config-ca-trustpoint)# `**`exit`**	Exits from Trustpoint Configuration mode.
Step 3	`hostname(config)# `**`crypto ca export`** `trustpoint` **`identity-certificate`** **`Example:`** `hostname(config)# crypto ca export local-asa identity-certificate`	Exports the certificate you created in Step 1. The certificate contents appear on the terminal screen. Copy the certificate from the terminal screen. This certificate enables Cisco UCM to validate the certificate that the ASA sends in the TLS handshake. On the local Cisco UCM, upload the certificate into the Cisco UCM trust store. See the Cisco Unified Communications Manager documentation for information. **Note** The subject name you enter while uploading the certificate to the local Cisco UCM is compared with the X.509 Subject Name field entered on the SIP Trunk Security Profile on Cisco UCM. For example, "Ent-local-domain-name" was entered in Step 1 of this task; therefore, "Ent-local-domain-name" should be entered in the Cisco UCM configuration.

Step 4	`hostname(config)# `**`crypto ca trustpoint`** `trustpoint_name`
	`hostname(config-ca-trustpoint)# `**`enroll terminal`**
	`Example:`
	`hostname(config)# crypto ca trustpoint local-ent-ucm`
	`hostname(config-ca-trustpoint)# enroll terminal`
		Creates a trustpoint for local Cisco UCM.
		
		Where `trustpoint_name` is the trustpoint for the local Cisco UCM.

Step 5	`hostname(config-ca-trustpoint)# `**`exit`**	Exits from Trustpoint Configuration mode.

	Commands	Purpose
Step 6	hostname(config)# **crypto ca authenticate** *trustpoint* **Example:** hostname(config)# crypto ca authenticate local-ent-ucm	Imports the certificate from local Cisco UCM. Where *trustpoint* is the trustpoint for the local Cisco UCM. Paste the certificate downloaded from the local Cisco UCM. This certificate enables the ASA to validate the certificate that Cisco UCM sends in the TLS handshake.
Step 7	hostname(config)# **tls-proxy** *proxy_name* hostname(config-tlsp)# **server trust-point** *proxy_trustpoint* hostname(config-tlsp)# **client trust-point** *proxy_trustpoint* hostname(config-tlsp)# **client cipher-suite** aes128-sha1 aes256-sha1 3des-sha1 null-sha1 **Example:** hostname(config)# tls-proxy local_to_remote-ent hostname(config-tlsp)# server trust-point local-ent-ucm hostname(config-tlsp)# client trust-point local-ent hostname(config-tlsp)# client cipher-suite aes128-sha1 aes256-sha1 3des-sha1 null-sha1	Updates the TLS proxy for **outbound** connections. Where *proxy_name* is the name you entered in Step 1 of the task Creating the TLS Proxy. Where *proxy_trustpoint* for the **server trust-point** command is the name you entered in Step 4 of this procedure. Where *proxy_trustpoint* for the **client trust-point** command is the name you entered in Step 2 of the task Creating Trustpoints and Generating Certificates. **Note** In this step, you are creating different trustpoints for the client and the server.
Step 8	hostname(config-tlsp)# **exit**	Exits from TLS Proxy Configuration mode.
Step 9	hostname(config)# **tls-proxy** proxy_name hostname(config-tlsp)# **server trust-point** *proxy_trustpoint* hostname(config-tlsp)# **client trust-point** *proxy_trustpoint* hostname(config-tlsp)# **client cipher-suite** aes128-sha1 aes256-sha1 3des-sha1 null-sha1 **Example:** hostname(config)# tls-proxy remote_to_local-ent hostname(config-tlsp)# server trust-point local-ent hostname(config-tlsp)# client trust-point local-ent-ucm hostname(config-tlsp)# client cipher-suite aes128-sha1 aes256-sha1 3des-sha1 null-sha1	Updates the TLS proxy for **inbound** connections. Where *proxy_name* is the name you entered in Step 5 of the task Creating the TLS Proxy. Where *proxy_trustpoint* for the **server trust-point** command is the name you entered in Step 2 of the task Creating Trustpoints and Generating Certificates. Where *proxy_trustpoint* for the **client trust-point** command is the name you entered in Step 4 of this procedure.
Step 10	hostname(config-tlsp)# **exit**	Exits from TLS Proxy Configuration mode.
Step 11	hostname(config)# **uc-ime** *uc_ime_name* hostname(config-uc-ime)# **ucm address** *ip_address* **trunk-security-mode secure** **Example:** hostname(config)# uc-ime local-ent-ime hostname(config-uc-ime)# ucm address 192.168.10.30 trunk-security-mode secure	Updates the Cisco Intercompany Media Engine Proxy for trunk-security-mode. Where *uc_ime_name* is the name you entered in Step 1 of the task Creating the Cisco Intercompany Media Engine Proxy. Only perform this step if you entered nonsecure in Step 3 of the task Creating the Cisco Intercompany Media Engine Proxy.

What to Do Next

Once you have configured the TLS within the enterprise, if necessary, configure off path signaling for an off path deployment. See (Optional) Configuring Off Path Signaling, page 19-31.

(Optional) Configuring Off Path Signaling

Perform this task only when you are configuring the Cisco Intercompany Media Engine Proxy as part of an off path deployment. You might choose to have an off path deployment when you want to use the Cisco Intercompany Media Engine but do not want to replace your existing Internet firewall with an ASA enabled with the Cisco Intercompany Media Engine Proxy.

In an off path deployment, the existing firewall that you have deployed in your environment is not capable of transmitting Cisco Intercompany Media Engine traffic.

Off path signaling requires that outside IP addresses translate to an inside IP address. The inside interface address can be used for this mapping service configuration. For the Cisco Intercompany Media Engine Proxy, the ASA creates dynamic mappings for external addresses to the internal IP address; therefore, using the dynamic NAT configuration on outbound calls, Cisco UCM sends SIP traffic to this internal IP address, and the ASA uses that mapping to determine the real destination on inbound calls. The static NAT or PAT mapping is used for inbound calls in an off path configuration.

Figure 19-9   Example for Configuring Off Path Signaling in an Off Path Deployment
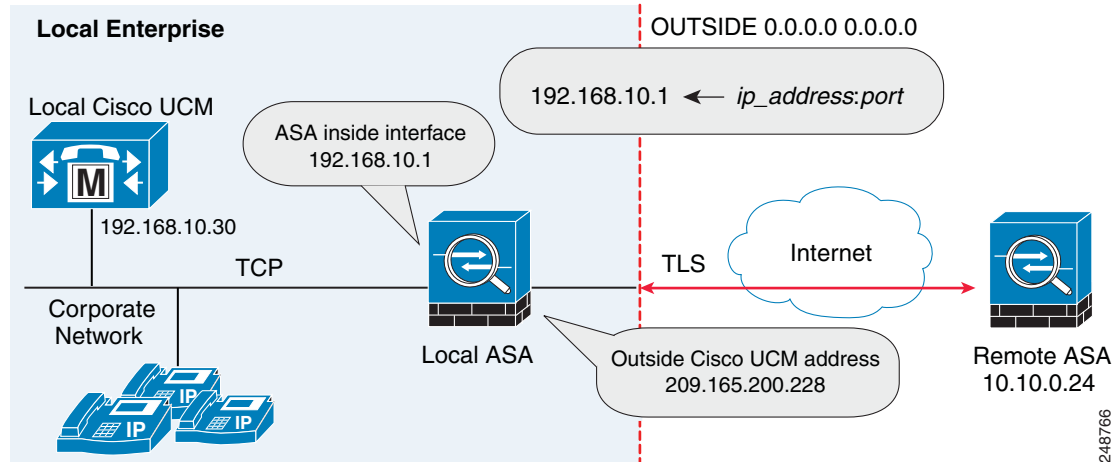
After you configure off path signaling, the ASA mapping service listens on interface "inside" for requests. When it receives a request, it creates a dynamic mapping for the "outside" as the destination interface.

To configure off path signaling for the Cisco Intercompany Media Engine Proxy, perform the following steps:

	Command	Purpose
Step 1	hostname(config)# **object network** *name* **Example:** hostname(config)# object network outside-any	For the off path ASA, creates a network object to represent all outside addresses.
Step 2	hostname(config-network-object)# **subnet** *ip_address* **Example:** hostname(config-network-object)# subnet 0.0.0.0 0.0.0.0	Specifies the IP address of the subnet.

	Command	Purpose
Step 3	`hostname(config-network-object)# nat (outside,inside) dynamic interface inside`	Creates a mapping for the Cisco UCM of remote enterprises.
Step 4	`hostname(config-network-object)# exit`	Exits from the objects configuration mode.

Step 5	`hostname(config)# uc-ime uc_ime_name`
	
	Example:
	`hostname(config)# uc-ime local-ent-ime`
		Specifies the Cisco Intercompany Media Engine Proxy that you created in the task Creating the Cisco Intercompany Media Engine Proxy, page 19-19.
		
		Where *uc_ime_name* is the name you specified in Step 1 of Creating the Cisco Intercompany Media Engine Proxy, page 19-19.

Step 6	`hostname(config)# mapping-service listening-interface interface_name [listening-port port] uc-ime-interface uc-ime-interface_name`
	
	Example:
	`hostname(config-uc-ime)# mapping-service listening-interface inside listening-port 8060 uc-ime-interface outside`
		For the off path ASA, adds the mapping service to the Cisco Intercompany Media Engine Proxy.
		
		Specifies the interface and listening port for the ASA mapping service.
		
		You can only configure one mapping server for the Cisco Intercompany Media Engine Proxy.
		
		Where *interface_name* is the name of the interface on which the ASA listens for the mapping requests.
		
		Where port is the TCP port on which the ASA listens for the mapping requests. The port number must be between 1024 and 65535 to avoid conflicts with other services on the device, such as Telnet or SSH. By default, the port number is TCP 8060.
		
		Where *uc-ime-interface_name* is the name of the interface that connects to the remote Cisco UCM.

This section contains the following sections:

Configuring the Cisco UC-IMC Proxy by using the UC-IME Proxy Pane

Use the Configure Cisco Intercompany Media Engine (UC-IME) proxy pane to add or edit a Cisco Intercompany Media Engine Proxy instance.

Note The Cisco Intercompany Media Engine Proxy does not appear as an option under the Unified Communications section of the navigation pane unless the license required for this proxy is installed on the ASA.

Use this pane to create the proxy instance; however, for the UC-IME proxy to be fully functionally, you must complete additional tasks, such as create the required NAT statements, ACLs, and MTA, set up the certificates, create the TLS Proxy, and enable SIP inspection.

Depending on whether the UC-IME proxy is deployed off path or in-line of Internet traffic, you must create the appropriate network objects with embedded NAT/PAT statements for the Cisco UCMs.

This pane is available from the Configuration > Firewall > Unified Communications > UC-IME Proxy.

Step 1   Open the Configuration > Firewall > Unified Communications > UC-IME Proxy pane.

Step 2   Check the Enable Cisco UC-IME proxy check box to enable the feature.

Step 3   In the Unified CM Servers area, enter an IP address or hostname for the Cisco Unified Communications Manager (Cisco UCM) or click the ellipsis to open a dialog and browse for an IP address or hostname.

Step 4   In the Trunk Security Mode field, click a security option. Specifying **secure** for Cisco UCM or Cisco UCM cluster indicates that Cisco UCM or Cisco UCM cluster is initiating TLS.

Step 5   Click **Add** to add the Cisco UCM for the Cisco Intercompany Media Engine Proxy. You must include an entry for each Cisco UCM in the cluster with Cisco Intercompany Media Engine that has a SIP trunk enabled.

Step 6   In the Ticket Epoch field, enter an integer from 1-255.

The epoch contains an integer that updates each time that the password is changed. When the proxy is configured the first time and a password entered for the first time, enter 1 for the epoch integer. Each time you change the password, increment the epoch to indicate the new password. You must increment the epoch value each time your change the password.

Typically, you increment the epoch sequentially; however, the ASA allows you to choose any value when you update the epoch.

If you change the epoch value, the current password is invalidated and you must enter a new password.

> ✎
> **Note**   The epoch and password that you configure in this step on the ASA must match the epoch and password that you configure on the Cisco Intercompany Media Engine server. See the Cisco Intercompany Media Engine server documentation for information.

Step 7   In the Ticket Password field, enter a minimum of 10 printable character from the US-ASCII character set. The allowed characters include 0x21 to 0x73 inclusive, and exclude the space character. The ticket password can be up to 64 characters. Confirm the password you entered. Only one password can be configured at a time.

Step 8   Check the Apply MTA to UC-IME Link proxy check box to associate the media termination address with the Cisco Intercompany Media Engine Proxy.

> ✎
> **Note**   You must create the media termination instance before you associate it with the Cisco Intercompany Media Engine Proxy. If necessary, click the Configure MTA button to configure a media termination address instance.

Step 9   If the Cisco Intercompany Media Engine Proxy is being configured as part of off path deployment, check the Enable off path address mapping service checkbox and configure the off path deployment settings:

   a.   From the Listening Interface field, select an ASA interface. This is the interface on which the ASA listens for the mapping requests.

   b.   In the Port field, enter a number between 1024 and 65535 as the TCP port on which the ASA listens for the mapping requests. The port number must be 1024 or higher to avoid conflicts with other services on the device, such as Telnet or SSH. By default, the port number is TCP 8060.

   c.   From the UC-IME Interface field, select an interface from the list. This is the interface that the ASA uses to connect to the remote Cisco UCM.

> **Note**   In an off path deployment any existing ASA that you have deployed in your environment are not capable of transmitting Cisco Intercompany Media Engine traffic. Off-path signaling requires that outside addresses are translated (using NAT) to an inside IP address. The inside interface address can be used for this mapping service configuration. For the Cisco Intercompany Media Engine Proxy, the ASA creates dynamic mappings for external addresses to the internal IP address.

Step 10   In the Fallback area, configure the fallback timer for the Cisco Intercompany Media Engine by specifying the following settings:

    a. In the Fallback Sensitivity File field, enter the path to a file in flash memory that the ASA uses for mid-call PSTN fallback. The file name that you enter must be the name of a file on disk that includes the .fbs file extension. Alternatively, click the Browse Flash button to locate and select the file from flash memory.

    b. In the Call Quality Evaluation Interval field, enter a number between 10-600 (in milliseconds). This number controls the frequency at which the ASA samples the RTP packets received from the Internet. The ASA uses the data sample to determine if fallback to the PSTN is needed for a call. By default, the length is 100 milliseconds for the timer.

    c. In the Notification Interval field, enter a number between 10-360 (in seconds). This number controls the amount of time that the ASA waits before notifying Cisco UCM whether to fall back to PSTN. By default, the length is 20 seconds for this timer.

> **Note**   When you change the fallback timer for the Cisco Intercompany Media Engine Proxy, ASDM automatically removes the proxy from SIP inspection and then reapplies SIP inspection when the proxy is re-enabled.

Step 11   Click Apply to save the configuration changes for the Cisco Intercompany Media Engine Proxy.

Configuring the Cisco UC-IMC Proxy by using the Unified Communications Wizard

To configure the Cisco Intercompany Media Engine Proxy by using ASDM, choose Wizards > Unified Communications Wizard from the menu. The Unified Communications Wizard opens. From the first page, select the Cisco Intercompany Media Engine Proxy option under the Business-to-Business section.

The wizard automatically creates the necessary TLS proxy, then guides you through creating the Intercompany Media Engine proxy, importing and installing the required certificates, and finally enables the SIP inspection for the Intercompany Media Engine traffic automatically.

The wizard guides you through these steps to create the Cisco Intercompany Media Engine Proxy:

Step 1   Select the Intercompany Media Engine Proxy option.

Step 2   Select the topology of the Cisco Intercompany Media Engine Proxy, namely whether the ASA is an edge firewall with all Internet traffic flowing through it or whether the ASA is off the path of the main Internet traffic (referred to as an off path deployment).

Step 3   Specify private network settings such as the Cisco UCM IP addresses and the ticket settings.

Step 4   Specify the public network settings.

Step 5   Specify the media termination address settings of Cisco UCM.

Step 6   Configure the local-side certificate management, namely the certificates that are exchanged between the local Cisco Unified Communications Manager servers and the ASA. The identity certificate that the wizard generates in this step needs to be installed on each Cisco Unified Communications Manager (UCM) server in the cluster with the proxy and each identity certificate from the Cisco UCMs need to be installed on the ASA. The certificates are used by the ASA and the Cisco UCMs to authenticate each other, respectively, during TLS handshakes. The wizard only supports self-signed certificates for this step.

Step 7   Configure the remote-side certificate management, namely the certificates that are exchanged between the remote server and the ASA. In this step, the wizard generates a certificate signing request (CSR). After successfully generating the identity certificate request for the proxy, the wizard prompts you to save the file.

You must send the CSR text file to a certificate authority (CA), for example, by pasting the text file into the CSR enrollment page on the CA website. When the CA returns the Identity Certificate, you must install it on the ASA. This certificate is presented to remote servers so that they can authenticate the ASA as a trusted server.

Finally, this step of the wizard assists you in installing the root certificates of the CA from the remote servers so that the ASA can determine that the remote servers are trusted.

The wizard completes by displaying a summary of the configuration created for Cisco Intercompany Media Engine. See the Unified Communications Wizard section in this documentation for more information.

This section describes how to certain options of the **show uc-ime** command to obtain troubleshooting information for the Cisco Intercompany Media Engine Proxy. See the command reference for detailed information about the syntax for these commands.

show uc-ime signaling-sessions

Displays the corresponding SIP signaling sessions stored by the Cisco Intercompany Media Engine Proxy. Use this command to troubleshoot media or signaling failure. The command also displays the fallback parameters extracted from the SIP message headers, whether RTP monitoring is enabled or disabled, and whether SRTP keys are set.

Through the use of the Cisco Intercompany Media Engine Proxy, not only signaling but also media is secured for communication. It provides signaling encryption and SRTP/RTP conversion with SRTP enforced on the Internet side. The Cisco Intercompany Media Engine Proxy inserts itself into the media path by modifying the SIP signaling messages from Cisco UCMs.The Cisco Intercompany Media Engine Proxy sits on the edge of the enterprise and inspects SIP signaling between SIP trunks created between enterprises. It terminates TLS signaling from the Internet and initiates TCP or TLS to the local Cisco UCM.

```
hostname# show uc-ime signaling-sessions
 1 in use, 3 most used
 inside 192.168.10.30:39608 outside 10.194.108.118:5070
   Local Media (audio) conn: 10.194.108.119/29824 to 10.194.108.109/21558
     Local SRTP key set : Remote SRTP key set
   Remote Media (audio) conn: 192.168.10.51/19520 to 192.168.10.3/30930
   Call-ID: ab6d7980-a7d11b08-50-1e0aa8c0@192.168.10.30
   FB Sensitivity: 3
   Session ID: 2948-32325449-0@81a985c9-f3a1-55a0-3b19-96549a027259
```

```
SIP Trunk URI: 81a985c9-f3a1-55a0-3b19-9654@UCM-30;maddr=192.168.10.30
Codec-name: G722
Payload type: 9
```

✎

Note   If calls are not going through the Cisco Intercompany Media Engine, you can also use the **show tls-proxy session** command to troubleshoot the success of the TLS handshake between the components in the Cisco Intercompany Media Engine system. See the command reference for information about this command.

show uc-ime signaling-sessions statistics

Displays statistical information about corresponding signaling sessions stored by Cisco Intercompany Media Engine Proxy. Failure of signaling sessions in the Cisco Intercompany Media Engine can occur for different call-related reasons; such as failure of ticket verification or domain name verification, or offering RTP over the Internet.

```
hostname# show uc-ime signaling-sessions statistics
10 in use, 20 most used
15 terminated
  Ticket integrity check failed: 2
  Ticket decode failed: 1
  Ticket epoch mismatch: 1
  Ticket DID mismatch: 0
  Ticket timestamp invalid: 4
  Ticket domain check failed: 2
  Ticket not found: 0
  Route domain name check failed: 1
  RTP over UC-IME: 2
```

✎

Note   Call-related failures, for example, can be due to the service policy rule being reconfigured or the primary ASA operating in failover mode. If a service policy rule for the Cisco Intercompany Media Engine Proxy is removed (by using the **no service policy** command) and reconfigured, the first call trasversing the ASA will fail. To resolve this issue, you must additionally enter the **clear connection** command and restart the ASA. If the failure is due to failover, the connections from the primary ASA are not synchronized to the standby ASA.

show uc-ime media-sessions detail

Displays the details about all active media sessions (calls) stored for the Cisco Intercompany Media Engine Proxy. Use this command to display output from successful calls. Additionally, use this command to troubleshoot problems with IP phone audio, such as one-way audio. If no calls are currently up, this output will be blank.

```
hostname(config)# show uc-ime media-sessions detail
 2 in use, 5 most used
 Media-session: 10.194.108.109/21558 :: client ip 192.168.10.51/19520
 Call ID: ab6d7980-a7d11b08-50-1e0aa8c0@192.168.10.30
 Session ID: 2948-32325449-0@81a985c9-f3a1-55a0-3b19-96549a027259
   Lcl SRTP conn 10.194.108.109/21558 to 10.194.108.119/29824 tx_pkts 20203 rx_pkts 20200
   refcnt 3 : created by Inspect SIP, passthrough not set
   RTP monitoring is enabled
      Failover_state              : 0
      Sum_all_packets             : 20196
      Codec_payload_format        : 9
      RTP_ptime_ms                : 20
      Max_RBLR_pct_x100           : 0
      Max_ITE_count_in_8_sec      : 0
```

```
        Max_BLS_ms                      :  0
        Max_PDV_usec                    :  1000
        Min_PDV_usec                    :  0
        Mov_avg_PDV_usec                :  109
        Total_ITE_count                 :  0
        Total_sec_count                 :  403
        Concealed_sec_count             :  0
        Severely_concealed_sec_count    :  0
        Max_call_interval_ms            :  118
        Total_SequenceNumber_Resets     :  0
 Media-session: 192.168.10.3/30930 :: client ip 10.194.108.119/29824
 Call ID: N/A
   Lcl RTP conn 192.168.10.3/30930 to 192.168.10.51/19520 tx_pkts 20201 rx_pkts 20203
```

show uc-ime fallback-notification statistics

Displays statistics about the PSTN fallback notifications to the Cisco UMC. Even if a call is sent over VoIP because the quality of the connection was good, the connection quality might worsen mid-call. To ensure an overall good experience for the end user, Cisco Intercompany Media Engine attempts to perform a mid-call fallback. Performing a mid-call fallback requires the adaptive security appliance to monitor the RTP packets coming from the Internet. If fallback is required, the adaptive security appliance sends a REFER message to Cisco UCM to tell it that it needs to fallback the call to PSTN.

Cisco Intercompany Media Engine uses a configurable hold-down timer to set the amount of time that adaptive security appliance waits before notifying Cisco UCM whether to fall back to PSTN.

```
hostname# show uc-ime fallback-notification statistics
UCM address: 172.23.32.37
 Total Notifications Sent: 10
```

show uc-ime mapping-service-sessions

When the Cisco Intercompany Media Engine Proxy is configured for an off path deployment, displays mapping-service requests and replies between the proxy and the local Cisco UMC. A TCP port on the ASA is configured to listen for mapping requests.

The port number must be 1024 or higher to avoid conflicts with other services on the device, such as Telnet or SSH. By default, the port number is TCP 8060.

```
Hostname# show uc-b2blink mapping-service-sessions
Total active sessions:  2
Session client (IP:Port)     Idle time
192.168.1.10:2001            0:01:01
192.168.1.20:3001            0:10:20
```

show uc-ime mapping-service-sessions statistics

Displays statistical information about the Cisco Intercompany Media Engine Proxy mapping service used in off path signaling.

```
Hostname# show uc-ime mapping-service-sessions statistics
Total active sessions:  2
Session client     Total       Responses   Failed      Pending     Idle
(IP:Port)          requests    sent        requests    responses   time
192.168.1.10:2001  10          9           1           0           0:01:01
192.168.1.20:3001  19          19          0           0           0:10:20
```

Feature History for Cisco Intercompany Media Engine Proxy

Table 19-1 lists the release history for this feature.

Table 19-1   Feature History for Cisco Phone Proxy

Feature Name	Releases	Feature Information
Cisco Intercompany Media Engine Proxy	8.3(1)	The Cisco Intercompany Media Engine Proxy was introduced.
		The following pane was added to the ASDM:
		Configuration > Firewall > Unified Communications > UC-IME Proxy
		The following wizard was added to ASDM, which allows you to configure the Unified Communication proxies (including the Cisco Intercompany Media Engine Proxy):
		Wizards > Unified Communications Wizard

P ART 5

Connection Settings and Quality of Service

Connection Settings

This chapter describes how to configure connection settings for connections that go through the ASA, or for management connections, that go to the ASA. Connection settings include:

- Maximum connections (TCP and UDP connections, embryonic connections, per-client connections)
- Connection timeouts
- Dead connection detection
- TCP sequence randomization
- TCP normalization customization
- TCP state bypass
- Global timeouts

This chapter includes the following sections:

Information About Connection Settings

This section describes why you might want to limit connections and includes the following topics:

TCP Intercept and Limiting Embryonic Connections

Limiting the number of embryonic connections protects you from a DoS attack. The ASA uses the per-client limits and the embryonic connection limit to trigger TCP Intercept, which protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. TCP Intercept uses the SYN cookies algorithm to prevent TCP SYN-flooding attacks. A SYN-flooding attack consists of a series of SYN packets usually originating from spoofed IP addresses. The constant flood of SYN packets keeps the server SYN queue full, which prevents it from servicing connection requests. When the embryonic connection threshold of a connection is crossed, the ASA acts as a proxy for the server and generates a SYN-ACK response to the client SYN request. When the ASA receives an ACK back from the client, it can then authenticate the client and allow the connection to the server.

Note   When you use TCP SYN cookie protection to protect servers from SYN attacks, you must set the embryonic connection limit lower than the TCP SYN backlog queue on the server that you want to protect. Otherwise, valid clients can nolonger access the server during a SYN attack.

To view TCP Intercept statistics, including the top 10 servers under attack, see Chapter 25, "Threat Detection."

Disabling TCP Intercept for Management Packets for Clientless SSL Compatibility

By default, TCP management connections have TCP Intercept always enabled. When TCP Intercept is enabled, it intercepts the 3-way TCP connection establishment handshake packets and thus deprives the ASA from processing the packets for clientless SSL. Clientless SSL requires the ability to process the 3-way handshake packets to provide selective ACK and other TCP options for clientless SSL connections. To disable TCP Intercept for management traffic, you can set the embryonic connection limit; only after the embryonic connection limit is reached is TCP Intercept enabled.

Dead Connection Detection (DCD)

DCD detects a dead connection and allows it to expire, without expiring connections that can still handle traffic. You configure DCD when you want idle, but valid connections to persist.

When you enable DCD, idle timeout behavior changes. With idle timeout, DCD probes are sent to each of the two end-hosts to determine the validity of the connection. If an end-host fails to respond after probes are sent at the configured intervals, the connection is freed, and reset values, if configured, are sent to each of the end-hosts. If both end-hosts respond that the connection is valid, the activity timeout is updated to the current time and the idle timeout is rescheduled accordingly.

Enabling DCD changes the behavior of idle-timeout handling in the TCP normalizer. DCD probing resets the idle timeout on the connections seen in the **show conn** command. To determine when a connection that has exceeded the configured timeout value in the timeout command but is kept alive due to DCD probing, the **show service-policy** command includes counters to show the amount of activity from DCD.

TCP Sequence Randomization

Each TCP connection has two ISNs: one generated by the client and one generated by the server. The ASA randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

TCP initial sequence number randomization can be disabled if required. For example:

*   If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.

*   If you use eBGP multi-hop through the ASA, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.

*   You use a WAAS device that requires the ASA not to randomize the sequence numbers of connections.

TCP Normalization

The TCP normalization feature identifies abnormal packets that the ASA can act on when they are detected; for example, the ASA can allow, drop, or clear the packets. TCP normalization helps protect the ASA from attacks. TCP normalization is always enabled, but you can customize how some features behave.

The TCP normalizer includes non-configurable actions and configurable actions. Typically, non-configurable actions that drop or clear connections apply to packets that are always bad. Configurable actions (as detailed in Customizing the TCP Normalizer with a TCP Map, page 20-6) might need to be customized depending on your network needs.

See the following guidelines for TCP normalization:

*   The normalizer does not protect from SYN floods. The ASA includes SYN flood protection in other ways.

*   The normalizer always sees the SYN packet as the first packet in a flow unless the ASA is in loose mode due to failover.

TCP State Bypass

By default, all traffic that goes through the ASA is inspected using the Adaptive Security Algorithm and is either allowed through or dropped based on the security policy. The ASA maximizes the firewall performance by checking the state of each packet (is this a new connection or an established

connection?) and assigning it to either the session management path (a new connection SYN packet), the fast path (an established connection), or the control plane path (advanced inspection). See the general operations configuration guide for more detailed information about the stateful firewall.

TCP packets that match existing connections in the fast path can pass through the ASA without rechecking every aspect of the security policy. This feature maximizes performance. However, the method of establishing the session in the fast path using the SYN packet, and the checks that occur in the fast path (such as TCP sequence number), can stand in the way of asymmetrical routing solutions: both the outbound and inbound flow of a connection must pass through the same ASA.

For example, a new connection goes to ASA 1. The SYN packet goes through the session management path, and an entry for the connection is added to the fast path table. If subsequent packets of this connection go through ASA 1, then the packets will match the entry in the fast path, and are passed through. But if subsequent packets go to ASA 2, where there was not a SYN packet that went through the session management path, then there is no entry in the fast path for the connection, and the packets are dropped. Figure 20-1 shows an asymmetric routing example where the outbound traffic goes through a different ASA than the inbound traffic:

Figure 20-1   Asymmetric Routing

If you have asymmetric routing configured on upstream routers, and traffic alternates between two ASAs, then you can configure TCP state bypass for specific traffic. TCP state bypass alters the way sessions are established in the fast path and disables the fast path checks. This feature treats TCP traffic much as it treats a UDP connection: when a non-SYN packet matching the specified networks enters the ASA, and there is not an fast path entry, then the packet goes through the session management path to establish the connection in the fast path. Once in the fast path, the traffic bypasses the fast path checks.

Licensing Requirements for Connection Settings

Model	License Requirement
ASAv	Standard or Premium License.
All other models	Base License.

Guidelines and Limitations

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent mode.

Failover Guidelines

Failover is supported.

TCP State Bypass Unsupported Features

The following features are not supported when you use TCP state bypass:

- Application inspection—Application inspection requires both inbound and outbound traffic to go through the same ASA, so application inspection is not supported with TCP state bypass.

- AAA authenticated sessions—When a user authenticates with one ASA, traffic returning via the other ASA will be denied because the user did not authenticate with that ASA.

- TCP Intercept, maximum embryonic connection limit, TCP sequence number randomization—The ASA does not keep track of the state of the connection, so these features are not applied.

- TCP normalization—The TCP normalizer is disabled.

- SSM and SSC functionality—You cannot use TCP state bypass and any application running on an SSM or SSC, such as IPS or CSC.

TCP State Bypass NAT Guidelines

Because the translation session is established separately for each ASA, be sure to configure static NAT on both ASAs for TCP state bypass traffic; if you use dynamic NAT, the address chosen for the session on ASA 1 will differ from the address chosen for the session on ASA 2.

Maximum Concurrent and Embryonic Connection Guidelines

Depending on the number of CPU cores on your ASA model, the maximum concurrent and embryonic connections may exceed the configured numbers due to the way each core manages connections. In the worst case scenario, the ASA allows up to n-1 extra connections and embryonic connections, where n is the number of cores. For example, if your model has 4 cores, if you configure 6 concurrent connections and 4 embryonic connections, you could have an additional 3 of each type. To determine the number of cores for your model, enter the **show cpu core** command.

Default Settings

TCP State Bypass

TCP state bypass is disabled by default.

Configuring Connection Settings

This section includes the following topics:

- Customizing the TCP Normalizer with a TCP Map, page 20-6
- Configuring Connection Settings, page 20-8
- Configuring Global Timeouts, page 20-9

Task Flow For Configuring Connection Settings

Step 1 For TCP normalization customization, create a TCP map according to the Customizing the TCP Normalizer with a TCP Map, page 20-6.

Step 2 For all connection settings except for global timeouts, configure a service policy according to Chapter 1, "Service Policy."

Step 3 Configure connection settings according to the Configuring Connection Settings, page 20-8.

Step 4 Configure global timeouts according to the Configuring Global Timeouts, page 20-9.

Customizing the TCP Normalizer with a TCP Map

To customize the TCP normalizer, first define the settings using a TCP map.

Detailed Steps

Step 1 Choose the **Configuration > Firewall > Objects > TCP Maps** pane, and click **Add**.

The Add TCP Map dialog box appears.

Step 2 In the TCP Map Name field, enter a name.

Step 3 In the Queue Limit field, enter the maximum number of out-of-order packets, between 0 and 250 packets.

The Queue Limit sets the maximum number of out-of-order packets that can be buffered and put in order for a TCP connection. The default is 0, which means this setting is disabled and the default system queue limit is used depending on the type of traffic:

- Connections for application inspection, IPS, and TCP check-retransmission have a queue limit of 3 packets. If the ASA receives a TCP packet with a different window size, then the queue limit is dynamically changed to match the advertised setting.
- For other TCP connections, out-of-order packets are passed through untouched.

If you set the Queue Limit to be 1 or above, then the number of out-of-order packets allowed for all TCP traffic matches this setting. For example, for application inspection, IPS, and TCP check-retransmission traffic, any advertised settings from TCP packets are ignored in favor of the Queue Limit setting. For other TCP traffic, out-of-order packets are now buffered and put in order instead of passed through untouched.

Step 4   In the Timeout field, set the maximum amount of time that out-of-order packets can remain in the buffer, between 1 and 20 seconds.

If they are not put in order and passed on within the timeout period, then they are dropped. The default is 4 seconds. You cannot change the timeout for any traffic if the Queue Limit is set to 0; you need to set the limit to be 1 or above for the Timeout to take effect.

Step 5   In the Reserved Bits area, click **Clear and allow**, **Allow only**, or **Drop**.

Allow only allows packets with the reserved bits in the TCP header.

Clear and allow clears the reserved bits in the TCP header and allows the packet.

Drop drops the packet with the reserved bits in the TCP header.

Step 6   Check any of the following options:

- Clear urgent flag—Clears the URG flag through the ASA. The URG flag is used to indicate that the packet contains information that is of higher priority than other data within the stream. The TCP RFC is vague about the exact interpretation of the URG flag, therefore end systems handle urgent offsets in different ways, which may make the end system vulnerable to attacks.

- Drop connection on window variation—Drops a connection that has changed its window size unexpectedly. The window size mechanism allows TCP to advertise a large window and to subsequently advertise a much smaller window without having accepted too much data. From the TCP specification, "shrinking the window" is strongly discouraged. When this condition is detected, the connection can be dropped.

- Drop packets that exceed maximum segment size—Drops packets that exceed MSS set by peer.

- Check if transmitted data is the same as original—Enables the retransmit data checks.

- Drop packets which have past-window sequence—Drops packets that have past-window sequence numbers, namely the sequence number of a received TCP packet is greater than the right edge of the TCP receiving window. If you do not check this option, then the Queue Limit must be set to 0 (disabled).

- Drop SYN Packets with data—Drops SYN packets with data.

- Enable TTL Evasion Protection—Enables the TTL evasion protection offered by the ASA. Do not enable this option if you want to prevent attacks that attempt to evade security policy.

- For example, an attacker can send a packet that passes policy with a very short TTL. When the TTL goes to zero, a router between the ASA and the endpoint drops the packet. It is at this point that the attacker can send a malicious packet with a long TTL that appears to the ASA to be a retransmission and is passed. To the endpoint host, however, it is the first packet that has been received by the attacker. In this case, an attacker is able to succeed without security preventing the attack.

- Verify TCP Checksum—Enables checksum verification.

- Drop SYNACK Packets with data—Drops TCP SYNACK packets that contain data.

- Drop packets with invalid ACK—Drops packets with an invalid ACK. You might see invalid ACKs in the following instances:

  - In the TCP connection SYN-ACK-received status, if the ACK number of a received TCP packet is not exactly same as the sequence number of the next TCP packet sending out, it is an invalid ACK.

- Whenever the ACK number of a received TCP packet is greater than the sequence number of the next TCP packet sending out, it is an invalid ACK.

✎

Note   TCP packets with an invalid ACK are automatically allowed for WAAS connections.

Step 7   To set TCP options, check any of the following options:

- Clear Selective Ack—Sets whether the selective-ack TCP option is allowed or cleared.
- Clear TCP Timestamp—Sets whether the TCP timestamp option is allowed or cleared.
- Clear Window Scale—Sets whether the window scale timestamp option is allowed or cleared.
- Range—Sets the valid TCP options ranges, which should fall within 6-7 and 9-255. The lower bound should be less than or equal to the upper bound. Choose **Allow** or **Drop** for each range.

Step 8   Click **OK**.

Configuring Connection Settings

To set connection settings, perform the following steps.

Detailed Steps

Step 1   Configure a service policy on the Configuration > Firewall > Service Policy Rules pane according to Chapter 1, "Service Policy."

You can configure connection limits as part of a new service policy rule, or you can edit an existing service policy.

Step 2   On the Rule Actions dialog box, click the **Connection Settings** tab.

Step 3   To set maximum connections, configure the following values in the Maximum Connections area:

- TCP & UDP Connections—Specifies the maximum number of simultaneous TCP and UDP connections for all clients in the traffic class, up to 2000000. The default is 0 for both protocols, which means the maximum possible connections are allowed.
- Embryonic Connections—Specifies the maximum number of embryonic connections per host up to 2000000. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. This limit enables the TCP Intercept feature. The default is **0**, which means the maximum embryonic connections. TCP Intercept protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. When the embryonic limit has been surpassed, the TCP intercept feature intercepts TCP SYN packets from clients to servers on a higher security level. SYN cookies are used during the validation process and help to minimize the amount of valid traffic being dropped. Thus, connection attempts from unreachable hosts will never reach the server.
- Per Client Connections—Specifies the maximum number of simultaneous TCP and UDP connections for each client up to 2000000. When a new connection is attempted by a client that already has opened the maximum per-client number of connections, the ASA rejects the connection and drops the packet.

> - Per Client Embryonic Connections—Specifies the maximum number of simultaneous TCP embryonic connections for each client up to 2000000. When a new TCP connection is requested by a client that already has the maximum per-client number of embryonic connections open through the ASA, the ASA proxies the request to the TCP Intercept feature, which prevents the connection.

Step 4   To configure connection timeouts, configure the following values in the TCP Timeout area:

> - Connection Timeout—Specifies the idle time until a connection slot (of *any* protocol, not just TCP) is freed. Enter 0:0:0 to disable timeout for the connection. This duration must be at least 5 minutes. The default is 1 hour.
>
> - Send reset to TCP endpoints before timeout—Specifies that the ASA should send a TCP reset message to the endpoints of the connection before freeing the connection slot.
>
> - Embryonic Connection Timeout—Specifies the idle time until an embryonic (half-open) connection slot is freed. Enter 0:0:0 to disable timeout for the connection. The default is 30 seconds.
>
> - Half Closed Connection Timeout—Sets the idle timeout period until a half-closed connection is closed, between 0:5:0 (for 9.1(1) and earlier) or 0:0:30 (for 9.1(2) and later) and 1193:0:0. The default is 0:10:0. Half-closed connections are not affected by DCD. Also, the ASA does not send a reset when taking down half-closed connections.

Step 5   To disable randomized sequence numbers, uncheck **Randomize Sequence Number**.

TCP initial sequence number randomization can be disabled if another in-line firewall is also randomizing the initial sequence numbers, because there is no need for both firewalls to be performing this action. However, leaving ISN randomization enabled on both firewalls does not affect the traffic.

Each TCP connection has two ISNs: one generated by the client and one generated by the server. The security appliance randomizes the ISN of the TCP SYN passing in the outbound direction. If the connection is between two interfaces with the same security level, then the ISN will be randomized in the SYN in both directions.

Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

Step 6   To configure TCP normalization, check **Use TCP Map**. Choose an existing TCP map from the drop-down list (if available), or add a new one by clicking **New**.

The Add TCP Map dialog box appears. See Customizing the TCP Normalizer with a TCP Map, page 20-6.

Step 7   Click **OK**.

Step 8   To set the time to live, check **Decrement time to live for a connection**.

Step 9   To enable TCP state bypass, in the Advanced Options area, check **TCP State Bypass**.

Step 10   Click **OK** or **Finish**.

Configuring Global Timeouts

The Configuration > Firewall > Advanced > Global Timeouts pane lets you set the timeout durations for use with the ASA. All durations are displayed in the format *hh*:*mm*:*ss*. It sets the idle time for the connection and translation slots of various protocols. If the slot has not been used for the idle time specified, the resource is returned to the free pool. TCP connection slots are freed approximately 60 seconds after a normal connection close sequence.

Fields

In all cases, except for Authentication absolute and Authentication inactivity, unchecking the check boxes means there is no timeout value. For those two cases, clearing the check box means to reauthenticate on every new connection.

- Connection—Modifies the idle time until a connection slot is freed. Enter 0:0:0 to disable timeout for the connection. This duration must be at least 5 minutes. The default is 1 hour.

- Half-closed—Modifies the idle time until a TCP half-closed connection closes. The minimum is 5 minutes. The default is 10 minutes. Enter 0:0:0 to disable timeout for a half-closed connection.

- UDP—Modifies the idle time until a UDP protocol connection closes. This duration must be at least 1 minute. The default is 2 minutes. Enter 0:0:0 to disable timeout.

- ICMP—Modifies the idle time after which general ICMP states are closed.

- H.323—Modifies the idle time until an H.323 media connection closes. The default is 5 minutes. Enter 0:0:0 to disable timeout.

- H.225—Modifies the idle time until an H.225 signaling connection closes. The H.225 default timeout is 1 hour (1:0:0). Setting the value of 0:0:0 means never close this connection. To close this connection immediately after all calls are cleared, a value of 1 second (0:0:1) is recommended.

- MGCP—Modifies the timeout value for MGCP which represents the idle time after which MGCP media ports are closed. The MGCP default timeout is 5 minutes (0:5:0). Enter 0:0:0 to disable timeout.

- MGCP PAT—Modifies the idle time after which an MGCP PAT translation is removed. The default is 5 minutes (0:5:0). The minimum time is 30 seconds. Uncheck the check box to return to the default value.

- TCP Proxy Reassembly—Configures the idle timeout after which buffered packets waiting for reassembly are dropped, between 0:0:10 and 1193:0:0. The default is 1 minute (0:1:0).

- Floating Connection—When multiple static routes exist to a network with different metrics, the ASA uses the one with the best metric at the time of connection creation. If a better route becomes available, then this timeout lets connections be closed so a connection can be reestablished to use the better route. The default is 0 (the connection never times out). To take advantage of this feature, change the timeout to a new value between 0:1:0 and 1193:0:0.

- SUNRPC—Modifies the idle time until a SunRPC slot is freed. This duration must be at least 1 minute. The default is 10 minutes. Enter 0:0:0 to disable timeout.

- SIP—Modifies the idle time until an SIP signalling port connection closes. This duration must be at least 5 minutes. The default is 30 minutes.

- SIP Media—Modifies the idle time until an SIP media port connection closes. This duration must be at least 1 minute. The default is 2 minutes.

- SIP Provisional Media—Modifies the timeout value for SIP provisional media connections, between 0:1:0 and 1193:0:0. The default is 2 minutes.

- SIP Invite—Modifies the idle time after which pinholes for PROVISIONAL responses and media xlates will be closed. The minimum value is 0:1:0, the maximum value is 0:30:0. The default value is 0:3:0.

- SIP Disconnect—Modifies the idle time after which SIP session is deleted if the 200 OK is not received for a CANCEL or a BYE message. The minimum value is 0:0:1, the maximum value is 0:10:0. The default value is 0:2:0.

- Authentication absolute—Modifies the duration until the authentication cache times out and you have to reauthenticate a new connection. This duration must be shorter than the Translation Slot value. The system waits until you start a new connection to prompt you again. Enter 0:0:0 to disable caching and reauthenticate on every new connection.

> **Note** Do not set this value to 0:0:0 if passive FTP is used on the connections.

> **Note** When Authentication Absolute = 0, HTTPS authentication may not work. If a browser initiates multiple TCP connections to load a web page after HTTPS authentication, the first connection is permitted through, but subsequent connections trigger authentication. As a result, users are continuously presented with an authentication page, even after successful authentication. To work around this, set the authentication absolute timeout to 1 second. This workaround opens a 1-second window of opportunity that might allow non-authenticated users to go through the firewall if they are coming from the same source IP address.

- Authentication inactivity—Modifies the idle time until the authentication cache times out and users have to reauthenticate a new connection. This duration must be shorter than the Translation Slot value.

- Translation Slot—Modifies the idle time until a translation slot is freed. This duration must be at least 1 minute. The default is 3 hours. Enter 0:0:0 to disable the timeout.

- (8.4(3) and later, not including 8.5(1) and 8.6(1)) PAT Translation Slot—Modifies the idle time until a PAT translation slot is freed, between 0:0:30 and 0:5:0. The default is 30 seconds. You may want to increase the timeout if upstream routers reject new connections using a freed PAT port because the previous connection might still be open on the upstream device.

Feature History for Connection Settings

Table 20-1 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 20-1   Feature History for Connection Settings

Feature Name	Platform Releases	Feature Information
TCP state bypass	8.2(1)	This feature was introduced. The following command was introduced: **set connection advanced-options tcp-state-bypass**.
Connection timeout for all protocols	8.2(2)	The idle timeout was changed to apply to all protocols, not just TCP. The following screen was modified: Configuration > Firewall > Service Policies > Rule Actions > Connection Settings.

Table 20-1   Feature History for Connection Settings (continued)

Feature Name	Platform Releases	Feature Information
Timeout for connections using a backup static route	8.2(5)/8.4(2)	When multiple static routes exist to a network with different metrics, the ASA uses the one with the best metric at the time of connection creation. If a better route becomes available, then this timeout lets connections be closed so a connection can be reestablished to use the better route. The default is 0 (the connection never times out). To take advantage of this feature, change the timeout to a new value. We modified the following screen: Configuration > Firewall > Advanced > Global Timeouts.
Configurable timeout for PAT xlate	8.4(3)	When a PAT xlate times out (by default after 30 seconds), and the ASA reuses the port for a new translation, some upstream routers might reject the new connection because the previous connection might still be open on the upstream device. The PAT xlate timeout is now configurable, to a value between 30 seconds and 5 minutes. We modified the following screen: Configuration > Firewall > Advanced > Global Timeouts. *This feature is not available in 8.5(1) or 8.6(1).*
Increased maximum connection limits for service policy rules	9.0(1)	The maximum number of connections for service policy rules was increased from 65535 to 2000000. We modified the following screen: Configuration > Firewall > Service Policy Rules > Connection Settings.
Decreased the half-closed timeout minimum value to 30 seconds	9.1(2)	The half-closed timeout minimum value for both the global timeout and connection timeout was lowered from 5 minutes to 30 seconds to provide better DoS protection. We modified the following screens: Configuration > Firewall > Service Policy Rules > Connection Settings Configuration > Firewall > Advanced > Global Timeouts.

CHAPTER **21**

Quality of Service

Have you ever participated in a long-distance phone call that involved a satellite connection? The conversation might be interrupted with brief, but perceptible, gaps at odd intervals. Those gaps are the time, called the latency, between the arrival of packets being transmitted over the network. Some network traffic, such as voice and video, cannot tolerate long latency times. Quality of service (QoS) is a feature that lets you give priority to critical traffic, prevent bandwidth hogging, and manage network bottlenecks to prevent packet drops.

Note For the ASASM, we suggest performing QoS on the switch instead of the ASASM. Switches have more capability in this area.

This chapter describes how to apply QoS policies and includes the following sections:

- Information About QoS, page 21-1
- Licensing Requirements for QoS, page 21-5
- Guidelines and Limitations, page 21-5
- Configuring QoS, page 21-6
- Monitoring QoS, page 21-11
- Feature History for QoS, page 21-14

Information About QoS

You should consider that in an ever-changing network environment, QoS is not a one-time deployment, but an ongoing, essential part of network design.

This section describes the QoS features supported by the ASA and includes the following topics:

- Supported QoS Features, page 21-2
- What is a Token Bucket?, page 21-2
- Information About Policing, page 21-3
- Information About Priority Queuing, page 21-3
- Information About Traffic Shaping, page 21-4
- DSCP and DiffServ Preservation, page 21-5

Supported QoS Features

The ASA supports the following QoS features:

- Policing—To prevent individual flows from hogging the network bandwidth, you can limit the maximum bandwidth used per flow. See Information About Policing, page 21-3 for more information.

- Priority queuing—For critical traffic that cannot tolerate latency, such as Voice over IP (VoIP), you can identify traffic for Low Latency Queuing (LLQ) so that it is always transmitted ahead of other traffic. See Information About Priority Queuing, page 21-3 for more information.

- Traffic shaping—If you have a device that transmits packets at a high speed, such as a ASA with Fast Ethernet, and it is connected to a low speed device such as a cable modem, then the cable modem is a bottleneck at which packets are frequently dropped. To manage networks with differing line speeds, you can configure the ASA to transmit packets at a fixed slower rate. See Information About Traffic Shaping, page 21-4 for more information.

What is a Token Bucket?

A token bucket is used to manage a device that regulates the data in a flow. For example, the regulator might be a traffic policer or a traffic shaper. A token bucket itself has no discard or priority policy. Rather, a token bucket discards tokens and leaves to the flow the problem of managing its transmission queue if the flow overdrives the regulator.

A token bucket is a formal definition of a rate of transfer. It has three components: a burst size, an average rate, and a time interval. Although the average rate is generally represented as bits per second, any two values may be derived from the third by the relation shown as follows:

average rate = burst size / time interval

Here are some definitions of these terms:

- Average rate—Also called the committed information rate (CIR), it specifies how much data can be sent or forwarded per unit time on average.

- Burst size—Also called the Committed Burst (Bc) size, it specifies in bits or bytes per burst how much traffic can be sent within a given unit of time to not create scheduling concerns. (For traffic shaping, it specifies bits per burst; for policing, it specifies bytes per burst.)

- Time interval—Also called the measurement interval, it specifies the time quantum in seconds per burst.

In the token bucket metaphor, tokens are put into the bucket at a certain rate. The bucket itself has a specified capacity. If the bucket fills to capacity, newly arriving tokens are discarded. Each token is permission for the source to send a certain number of bits into the network. To send a packet, the regulator must remove from the bucket a number of tokens equal in representation to the packet size.

If not enough tokens are in the bucket to send a packet, the packet either waits until the bucket has enough tokens (in the case of traffic shaping) or the packet is discarded or marked down (in the case of policing). If the bucket is already full of tokens, incoming tokens overflow and are not available to future packets. Thus, at any time, the largest burst a source can send into the network is roughly proportional to the size of the bucket.

Note that the token bucket mechanism used for traffic shaping has both a token bucket and a data buffer, or queue; if it did not have a data buffer, it would be a policer. For traffic shaping, packets that arrive that cannot be sent immediately are delayed in the data buffer.

For traffic shaping, a token bucket permits burstiness but bounds it. It guarantees that the burstiness is bounded so that the flow will never send faster than the token bucket capacity, divided by the time interval, plus the established rate at which tokens are placed in the token bucket. See the following formula:

(token bucket capacity in bits / time interval in seconds) + established rate in bps = maximum flow speed in bps

This method of bounding burstiness also guarantees that the long-term transmission rate will not exceed the established rate at which tokens are placed in the bucket.

Information About Policing

Policing is a way of ensuring that no traffic exceeds the maximum rate (in bits/second) that you configure, thus ensuring that no one traffic flow or class can take over the entire resource. When traffic exceeds the maximum rate, the ASA drops the excess traffic. Policing also sets the largest single burst of traffic allowed.

Information About Priority Queuing

LLQ priority queuing lets you prioritize certain traffic flows (such as latency-sensitive traffic like voice and video) ahead of other traffic.

The ASA supports two types of priority queuing:

- Standard priority queuing—Standard priority queuing uses an LLQ priority queue on an interface (see Configuring the Standard Priority Queue for an Interface, page 21-8), while all other traffic goes into the "best effort" queue. Because queues are not of infinite size, they can fill and overflow. When a queue is full, any additional packets cannot get into the queue and are dropped. This is called *tail drop*. To avoid having the queue fill up, you can increase the queue buffer size. You can also fine-tune the maximum number of packets allowed into the transmit queue. These options let you control the latency and robustness of the priority queuing. Packets in the LLQ queue are always transmitted before packets in the best effort queue.

- Hierarchical priority queuing—Hierarchical priority queuing is used on interfaces on which you enable a traffic shaping queue. A subset of the shaped traffic can be prioritized. The standard priority queue is not used. See the following guidelines about hierarchical priority queuing:

  - Priority packets are always queued at the head of the shape queue so they are always transmitted ahead of other non-priority queued packets.

  - Priority packets are never dropped from the shape queue unless the sustained rate of priority traffic exceeds the shape rate.

  - For IPsec-encrypted packets, you can only match traffic based on the DSCP or precedence setting.

  - IPsec-over-TCP is not supported for priority traffic classification.

Information About Traffic Shaping

Traffic shaping is used to match device and link speeds, thereby controlling packet loss, variable delay, and link saturation, which can cause jitter and delay.

Note Traffic shaping is only supported on the ASA 5505.

- Traffic shaping must be applied to all outgoing traffic on a physical interface or in the case of the ASA 5505, on a VLAN. You cannot configure traffic shaping for specific types of traffic.

- Traffic shaping is implemented when packets are ready to be transmitted on an interface, so the rate calculation is performed based on the actual size of a packet to be transmitted, including all the possible overhead such as the IPsec header and L2 header.

- The shaped traffic includes both through-the-box and from-the-box traffic.

- The shape rate calculation is based on the standard token bucket algorithm. The token bucket size is twice the Burst Size value. See What is a Token Bucket?, page 21-2.

- When bursty traffic exceeds the specified shape rate, packets are queued and transmitted later. Following are some characteristics regarding the shape queue (for information about hierarchical priority queuing, see Information About Priority Queuing, page 21-3):

  – The queue size is calculated based on the shape rate. The queue can hold the equivalent of 200-milliseconds worth of shape rate traffic, assuming a 1500-byte packet. The minimum queue size is 64.

  – When the queue limit is reached, packets are tail-dropped.

  – Certain critical keep-alive packets such as OSPF Hello packets are never dropped.

  – The time interval is derived by $time\_interval = burst\_size / average\_rate$. The larger the time interval is, the burstier the shaped traffic might be, and the longer the link might be idle. The effect can be best understood using the following exaggerated example:

  Average Rate = 1000000

  Burst Size = 1000000

  In the above example, the time interval is 1 second, which means, 1 Mbps of traffic can be bursted out within the first 10 milliseconds of the 1-second interval on a 100 Mbps FE link and leave the remaining 990 milliseconds idle without being able to send any packets until the next time interval. So if there is delay-sensitive traffic such as voice traffic, the Burst Size should be reduced compared to the average rate so the time interval is reduced.

How QoS Features Interact

You can configure each of the QoS features alone if desired for the ASA. Often, though, you configure multiple QoS features on the ASA so you can prioritize some traffic, for example, and prevent other traffic from causing bandwidth problems.

See the following supported feature combinations per interface:

- Standard priority queuing (for specific traffic) + Policing (for the rest of the traffic).

  You cannot configure priority queuing and policing for the same set of traffic.

- Traffic shaping (for all traffic on an interface) + Hierarchical priority queuing (for a subset of traffic).

You cannot configure traffic shaping and standard priority queuing for the same interface; only hierarchical priority queuing is allowed. For example, if you configure standard priority queuing for the global policy, and then configure traffic shaping for a specific interface, the feature you configured last is rejected because the global policy overlaps the interface policy.

Typically, if you enable traffic shaping, you do not also enable policing for the same traffic, although the ASA does not restrict you from configuring this.

DSCP and DiffServ Preservation

- DSCP markings are preserved on all traffic passing through the ASA.
- The ASA does not locally mark/remark any classified traffic, but it honors the Expedited Forwarding (EF) DSCP bits of every packet to determine if it requires "priority" handling and will direct those packets to the LLQ.
- DiffServ marking is preserved on packets when they traverse the service provider backbone so that QoS can be applied in transit (QoS tunnel pre-classification).

Licensing Requirements for QoS

The following table shows the licensing requirements for this feature:

Model	License Requirement
ASAv	Standard or Premium License.
All other models	Base License.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single context mode only. Does not support multiple context mode.

Firewall Mode Guidelines

Supported in routed firewall mode only. Does not support transparent firewall mode.

IPv6 Guidelines

Does not support IPv6.

Model Guidelines

- Traffic shaping is only supported on the ASA 5505. Multi-core models (such as the ASA 5500-X) do not support shaping.
- (ASA 5512-X through ASA 5555-X) Priority queuing is not supported on the Management 0/0 interface.

- (ASASM) Only policing is supported.

Additional Guidelines and Limitations

- QoS is applied unidirectionally; only traffic that enters (or exits, depending on the QoS feature) the interface to which you apply the policy map is affected. See Feature Directionality, page 1-2 for more information.

- For traffic shaping, you can only use the **class-default** class map, which is automatically created by the ASA, and which matches all traffic.

- For priority traffic, you cannot use the **class-default** class map.

- For hierarchical priority queuing, for encrypted VPN traffic, you can only match traffic based on the DSCP or precedence setting; you cannot match a tunnel group.

- For hierarchical priority queuing, IPsec-over-TCP traffic is not supported.

- You cannot configure traffic shaping and standard priority queuing for the same interface; only hierarchical priority queuing is allowed.

- For standard priority queuing, the queue must be configured for a physical interface or, for the ASA 5505 or ASASM, a VLAN.

- For policing, to-the-box traffic is not supported.

- For policing, traffic to and from a VPN tunnel bypass interface is not supported.

- For policing, when you match a tunnel group class map, only outbound policing is supported.

Configuring QoS

This section includes the following topics:

Determining the Queue and TX Ring Limits for a Standard Priority Queue

To determine the priority queue and TX ring limits, use the worksheets below.

Table 21-1 shows how to calculate the priority queue size. Because queues are not of infinite size, they can fill and overflow. When a queue is full, any additional packets cannot get into the queue and are dropped (called *tail drop*). To avoid having the queue fill up, you can adjust the queue buffer size according to the Configuring the Standard Priority Queue for an Interface, page 21-8.

Table 21-1   Queue Limit Worksheet

Step 1						
_____ *Outbound bandwidth (Mbps or Kbps)[1]*	Mbps	**x** **125**		**=**	_____ *# of bytes/ms*	
	Kbps	**x** **.125**		**=**	_____ *# of bytes/ms*	
Step 2						
_____ *# of bytes/ms from Step 1*	**÷**	_____ *Average packet size (bytes)[2]*	**x**	_____ *Delay (ms)[3]*	**=**	_____ *Queue limit (# of packets)*

1. For example, DSL might have an uplink speed of 768 Kbps. Check with your provider.

2. Determine this value from a codec or sampling size. For example, for VoIP over VPN, you might use 160 bytes. We recommend 256 bytes if you do not know what size to use.

3. The delay depends on your application. For example, the recommended maximum delay for VoIP is 200 ms. We recommend 500 ms if you do not know what delay to use.

Table 21-2 shows how to calculate the TX ring limit. This limit determines the maximum number of packets allowed into the Ethernet transmit driver before the driver pushes back to the queues on the interface to let them buffer packets until the congestion clears. This setting guarantees that the hardware-based transmit ring imposes a limited amount of extra latency for a high-priority packet.

Table 21-2   TX Ring Limit Worksheet

Step 1						
_____ *Outbound bandwidth (Mbps or Kbps)[1]*	Mbps	**x** **125**		**=**	_____ *# of bytes/ms*	
	Kbps	**x** **0.125**		**=**	_____ *# of bytes/ms*	
Step 2						
_____ *# of bytes/ms from Step 1*	**÷**	_____ *Maximum packet size (bytes)[2]*	**x**	_____ *Delay (ms)[3]*	**=**	_____ *TX ring limit (# of packets)*

1. For example, DSL might have an uplink speed of 768 Kbps. Check with your provider.

2. Typically, the maximum size is 1538 bytes, or 1542 bytes for tagged Ethernet. If you allow jumbo frames (if supported for your platform), then the packet size might be larger.

3. The delay depends on your application. For example, to control jitter for VoIP, you should use 20 ms.

Configuring the Standard Priority Queue for an Interface

If you enable standard priority queuing for traffic on a physical interface, then you need to also create the priority queue on each interface. Each physical interface uses two queues: one for priority traffic, and the other for all other traffic. For the other traffic, you can optionally configure policing.

Note   The standard priority queue is not required for hierarchical priority queuing with traffic shaping; see Information About Priority Queuing, page 21-3 for more information.

Restrictions

- (ASASM) The ASASM does not support priority queuing.
- (ASA 5512-X through ASA 5555-X) Priority queuing is not supported on the Management 0/0 interface.

Detailed Steps

Step 1   Go to Configuration > Device Management > Advanced > Priority Queue, and click **Add**.

The Add Priority Queue dialog box displays.

Step 2   From the Interface drop-down list, choose the physical interface name on which you want to enable the priority queue, or for the ASA 5505 or ASASM, the VLAN interface name.

Step 3   To change the size of the priority queues, in the Queue Limit field, enter the number of average, 256-byte packets that the specified interface can transmit in a 500-ms interval.

A packet that stays more than 500 ms in a network node might trigger a timeout in the end-to-end application. Such a packet can be discarded in each network node.

Because queues are not of infinite size, they can fill and overflow. When a queue is full, any additional packets cannot get into the queue and are dropped (called *tail drop*). To avoid having the queue fill up, you can use this option to increase the queue buffer size.

The upper limit of the range of values for this option is determined dynamically at run time. The key determinants are the memory needed to support the queues and the memory available on the device.

The Queue Limit that you specify affects both the higher priority low-latency queue and the best effort queue.

Step 4   To specify the depth of the priority queues, in the Transmission Ring Limit field, enter the number of maximum 1550-byte packets that the specified interface can transmit in a 10-ms interval.

This setting guarantees that the hardware-based transmit ring imposes no more than 10-ms of extra latency for a high-priority packet.

This option sets the maximum number of low-latency or normal priority packets allowed into the Ethernet transmit driver before the driver pushes back to the queues on the interface to let them buffer packets until the congestion clears.

The upper limit of the range of values is determined dynamically at run time. The key determinants are the memory needed to support the queues and the memory available on the device.

The Transmission Ring Limit that you specify affects both the higher priority low-latency queue and the best-effort queue.

Configuring a Service Rule for Standard Priority Queuing and Policing

You can configure standard priority queuing and policing for different class maps within the same policy map. See How QoS Features Interact, page 21-4 for information about valid QoS configurations.

To create a policy map, perform the following steps.

Restrictions

- You cannot use the **class-default** class map for priority traffic.
- You cannot configure traffic shaping and standard priority queuing for the same interface; only hierarchical priority queuing is allowed.
- (ASASM) The ASASM only supports policing.
- For policing, to-the-box traffic is not supported.
- For policing, traffic to and from a VPN tunnel bypass interface is not supported.
- For policing, when you match a tunnel group class map, only outbound policing is supported.

Guidelines

- For priority traffic, identify only latency-sensitive traffic.
- For policing traffic, you can choose to police all other traffic, or you can limit the traffic to certain types.

Detailed Steps

Step 1   To configure priority queuing, configure a service policy rule in the Configuration > Firewall > Service Policy Rules pane according to Chapter 1, "Service Policy."

You can configure QoS as part of a new service policy rule, or you can edit an existing service policy.

Step 2   In the Rule Actions dialog box, click the **QoS** tab.

Step 3   Click **Enable priority for this flow**.

If this service policy rule is for an individual interface, ASDM automatically creates the priority queue for the interface (Configuration > Device Management > Advanced > Priority Queue; for more information, see Configuring the Standard Priority Queue for an Interface, page 21-8). If this rule is for the global policy, then you need to manually add the priority queue to one or more interfaces *before* you configure the service policy rule.

Step 4   Click **Finish**. The service policy rule is added to the rule table.

Step 5   To configure policing, configure a service policy rule for the same interface in the Configuration > Firewall > Service Policy Rules pane according to Chapter 1, "Service Policy."

For policing traffic, you can choose to police all traffic that you are not prioritizing, or you can limit the traffic to certain types.

Step 6 In the Rule Actions dialog box, click the **QoS** tab.

Step 7 Click **Enable policing**, then check the **Input policing** or **Output policing** (or both) check boxes to enable the specified type of traffic policing. For each type of traffic policing, configure the following fields:

- Committed Rate—The rate limit for this traffic flow; this is a value in the range 8000-2000000000, specifying the maximum speed (bits per second) allowed.

- Conform Action—The action to take when the rate is less than the conform-burst value. Values are transmit or drop.

- Exceed Action—Take this action when the rate is between the conform-rate value and the conform-burst value. Values are transmit or drop.

- Burst Rate—A value in the range 1000-512000000, specifying the maximum number of instantaneous bytes allowed in a sustained burst before throttling to the conforming rate value.

Step 8 Click **Finish**. The service policy rule is added to the rule table.

Step 9 Click **Apply** to send the configuration to the device.

Configuring a Service Rule for Traffic Shaping and Hierarchical Priority Queuing

You can configure traffic shaping for all traffic on an interface, and optionally hierarchical priority queuing for a subset of latency-sensitive traffic.

Guidelines

- One side-effect of priority queuing is packet re-ordering. For IPsec packets, out-of-order packets that are not within the anti-replay window generate warning syslog messages. These warnings are false alarms in the case of priority queuing. You can configure the IPsec anti-replay window size to avoid possible false alarms. See the Configuration > VPN > IPsec > IPsec Rules > Enable Anti-replay window size option in the VPN configuration guide.

- For hierarchical priority queuing, you do not need to create a priority queue on an interface.

Restrictions

- For hierarchical priority queuing, for encrypted VPN traffic, you can only match traffic based on the DSCP or precedence setting; you cannot match a tunnel group.

- For hierarchical priority queuing, IPsec-over-TCP traffic is not supported.

- Traffic shaping is only supported on the ASA 5505. Multi-core models (such as the ASA 5500-X) do not support shaping.

- For traffic shaping, you can only use the **class-default** class map, which is automatically created by the ASA, and which matches all traffic.

- You cannot configure traffic shaping and standard priority queuing for the same interface; only hierarchical priority queuing is allowed. See How QoS Features Interact, page 21-4 for information about valid QoS configurations.

- You cannot configure traffic shaping in the global policy.

Detailed Steps

Step 1 Configure a service policy on the Configuration > Firewall > Service Policy Rules pane according to Chapter 1, "Service Policy."

You can configure QoS as part of a new service policy rule, or you can edit an existing service policy.

Step 2 In the Rule Actions dialog box, click the **QoS** tab.

Step 3 Click **Enable traffic shaping**, and configure the following fields:

- Average Rate—Sets the average rate of traffic in bits per second over a given fixed time period, between 64000 and 154400000. Specify a value that is a multiple of 8000.

- Burst Size—Sets the average burst size in bits that can be transmitted over a given fixed time period, between 2048 and 154400000. Specify a value that is a multiple of 128. If you do not specify the Burst Size, the default value is equivalent to 4-milliseconds of traffic at the specified Average Rate. For example, if the average rate is 1000000 bits per second, 4 ms worth = 1000000 * 4/1000 = 4000.

Step 4 (Optional) To configure priority queuing for a subset of shaped traffic:

  a. Click **Enforce priority to selected shape traffic**.

  b. Click **Configure** to identify the traffic that you want to prioritize.

    You are prompted to identify the traffic for which you want to apply priority queuing.

  c. After you identify the traffic (see Adding a Service Policy Rule for Through Traffic, page 1-8), click **Next**.

  d. Click **Enable priority for this flow**.

  e. Click **Finish**.

    You return to the QoS tab.

Step 5 Click **Finish**. The service policy rule is added to the rule table.

Step 6 Click **Apply** to send the configuration to the device.

Monitoring QoS

To monitor QoS in ASDM, you can enter commands at the Command Line Interface tool. This section includes the following topics:

- Viewing QoS Police Statistics, page 21-12
- Viewing QoS Standard Priority Statistics, page 21-12
- Viewing QoS Shaping Statistics, page 21-12
- Viewing QoS Standard Priority Queue Statistics, page 21-13

Viewing QoS Police Statistics

To view the QoS statistics for traffic policing, use the **show service-policy** command with the **police** keyword:

```
hostname# show service-policy police
```

The following is sample output for the **show service-policy police** command:

```
hostname# show service-policy police

Global policy:
    Service-policy: global_fw_policy

Interface outside:
    Service-policy: qos
        Class-map: browse
            police Interface outside:
                cir 56000 bps, bc 10500 bytes
                conformed 10065 packets, 12621510 bytes; actions: transmit
                exceeded 499 packets, 625146 bytes; actions: drop
                conformed 5600 bps, exceed 5016 bps
        Class-map: cmap2
            police Interface outside:
                cir 200000 bps, bc 37500 bytes
                conformed 17179 packets, 20614800 bytes; actions: transmit
                exceeded 617 packets, 770718 bytes; actions: drop
                conformed 198785 bps, exceed 2303 bps
```

Viewing QoS Standard Priority Statistics

To view statistics for service policies implementing the **priority** command, use the **show service-policy** command with the **priority** keyword:

```
hostname# show service-policy priority
```

The following is sample output for the **show service-policy priority** command:

```
hostname# show service-policy priority
Global policy:
    Service-policy: global_fw_policy
Interface outside:
    Service-policy: qos
        Class-map: TG1-voice
            Priority:
                Interface outside: aggregate drop 0, aggregate transmit 9383
```

Note   "Aggregate drop" denotes the aggregated drop in this interface; "aggregate transmit" denotes the aggregated number of transmitted packets in this interface.

Viewing QoS Shaping Statistics

To view statistics for service policies implementing the **shape** command, use the **show service-policy** command with the **shape** keyword:

```
hostname# show service-policy shape
```

The following is sample output for the **show service-policy shape** command:

```
hostname# show service-policy shape
Interface outside
  Service-policy: shape
    Class-map: class-default

      Queueing
      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 0/0

      shape (average) cir 2000000, bc 8000, be 8000
```

The following is sample output of the **show service policy shape** command, which includes service policies that include the **shape** command and the **service-policy** command that calls the hierarchical priority policy and the related statistics:

```
hostname# show service-policy shape

Interface outside:
  Service-policy: shape
    Class-map: class-default

      Queueing
      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 0/0

      shape (average) cir 2000000, bc 16000, be 16000

      Service-policy: voip
        Class-map: voip

          Queueing
          queue limit 64 packets
          (queue depth/total drops/no-buffer drops) 0/0/0
          (pkts output/bytes output) 0/0
        Class-map: class-default

          queue limit 64 packets
          (queue depth/total drops/no-buffer drops) 0/0/0
          (pkts output/bytes output) 0/0
```

Viewing QoS Standard Priority Queue Statistics

To display the priority-queue statistics for an interface, use the **show priority-queue statistics** command in privileged EXEC mode. The results show the statistics for both the best-effort (BE) queue and the low-latency queue (LLQ). The following example shows the use of the **show priority-queue statistics** command for the interface named test, and the command output.

```
hostname# show priority-queue statistics test

Priority-Queue Statistics interface test

Queue Type        = BE
Packets Dropped   = 0
Packets Transmit  = 0
Packets Enqueued  = 0
```

```
                    Current Q Length   = 0
                    Max Q Length       = 0

                    Queue Type         = LLQ
                    Packets Dropped    = 0
                    Packets Transmit   = 0
                    Packets Enqueued   = 0
                    Current Q Length   = 0
                    Max Q Length       = 0
                    hostname#
```

In this statistical report, the meaning of the line items is as follows:

- "Packets Dropped" denotes the overall number of packets that have been dropped in this queue.

- "Packets Transmit" denotes the overall number of packets that have been transmitted in this queue.

- "Packets Enqueued" denotes the overall number of packets that have been queued in this queue.

- "Current Q Length" denotes the current depth of this queue.

- "Max Q Length" denotes the maximum depth that ever occurred in this queue.

Feature History for QoS

Table 21-3 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 21-3   *Feature History for QoS*

Feature Name	Platform Releases	Feature Information
Priority queuing and policing	7.0(1)	We introduced QoS priority queuing and policing. We introduced the following screens: Configuration > Device Management > Advanced > Priority Queue Configuration > Firewall > Service Policy Rules
Shaping and hierarchical priority queuing	7.2(4)/8.0(4)	We introduced QoS shaping and hierarchical priority queuing. We modified the following screen: Configuration > Firewall > Service Policy Rules.
Ten Gigabit Ethernet support for a standard priority queue on the ASA 5585-X	8.2(3)/8.4(1)	We added support for a standard priority queue on Ten Gigabit Ethernet interfaces for the ASA 5585-X.

CHAPTER **22**

Troubleshooting Connections and Resources

This chapter describes how to troubleshoot the ASA and includes the following sections:

- Testing Your Configuration, page 22-1
- Monitoring Performance, page 22-8
- Monitoring System Resources, page 22-9
- Monitoring Connections, page 22-11
- Monitoring Per-Process CPU Usage, page 22-12

Testing Your Configuration

This section describes how to test connectivity for the single mode ASA or for each security context, how to ping the ASA interfaces, and how to allow hosts on one interface to ping through to hosts on another interface.

This section includes the following topics:

- Pinging ASA Interfaces, page 22-1
- Verifying ASA Configuration and Operation, and Testing Interfaces Using Ping, page 22-3
- Determining Packet Routing with Traceroute, page 22-6
- Tracing Packets with Packet Tracer, page 22-7

Pinging ASA Interfaces

To test whether the ASA interfaces are up and running and that the ASA and connected routers are operating correctly, you can ping the ASA interfaces.

To ping the ASA interfaces, perform the following steps:

Step 1 Draw a diagram of your single-mode ASA or security context that shows the interface names, security levels, and IP addresses.

✎

Note Although this procedure uses IP addresses, the **ping** command also supports DNS names and names that are assigned to a local IP address with the **name** command.

The diagram should also include any directly connected routers and a host on the other side of the router from which you will ping the ASA. (See Figure 22-1.)

Figure 22-1   Network Diagram with Interfaces, Routers, and Hosts

Step 2   Ping each ASA interface from the directly connected routers. For transparent mode, ping the management IP address. This test ensures that the ASA interfaces are active and that the interface configuration is correct.

A ping might fail if the ASA interface is not active, the interface configuration is incorrect, or if a switch between the ASA and a router is down (see Figure 22-2). In this case, no debugging messages or syslog messages appear, because the packet never reaches the ASA.

Figure 22-2   Ping Failure at the ASA Interface

If the ping reaches the ASA, and it responds, debugging messages similar to the following appear:

```
ICMP echo reply (len 32 id 1 seq 256) 209.165.201.1 > 209.165.201.2
ICMP echo request (len 32 id 1 seq 512) 209.165.201.2 > 209.165.201.1
```

If the ping reply does not return to the router, then a switch loop or redundant IP addresses may exist (see Figure 22-3).

Figure 22-3   *Ping Failure Because of IP Addressing Problems*

Step 3   Ping each ASA interface from a remote host. For transparent mode, ping the management IP address. This test checks whether the directly connected router can route the packet between the host and the ASA, and whether the ASA can correctly route the packet back to the host.

A ping might fail if the ASA does not have a return route to the host through the intermediate router (see Figure 22-4). In this case, the debugging messages show that the ping was successful, but syslog message 110001 appears, indicating a routing failure has occurred.

Figure 22-4   *Ping Failure Because the ASA Has No Return Route*

Verifying ASA Configuration and Operation, and Testing Interfaces Using Ping

The Ping tool is useful for verifying the configuration and operation of the ASA and surrounding communications links, as well as for testing other network devices.

This section includes the following topics:

- Information About Ping, page 22-3
- Pinging From an ASA Interface, page 22-4
- Pinging to an ASA Interface, page 22-4
- Pinging Through the ASA Interface, page 22-4
- Troubleshooting the Ping Tool, page 22-4
- Using the Ping Tool, page 22-5

Information About Ping

A ping is sent to an IP address and it returns a reply. This process enables network devices to discover, identify, and test each other.

The Ping tool uses ICMP (as described in RFC 777 and RFC 792) to define an echo request-and-reply transaction between two network devices. The echo request packet is sent to the IP address of a network device. The receiving device reverses the source and destination address and sends the packet back as the echo reply.

Administrators can use the ASDM Ping interactive diagnostic tool in these ways:

- Loopback testing of two interfaces—A ping may be initiated from one interface to another on the same ASA, as an external loopback test to verify basic "up" status and operation of each interface.

- Pinging to an ASA—The Ping tool can ping an interface on another ASA to verify that it is up and responding.

- Pinging through an ASA—Ping packets originating from the Ping tool may pass through an intermediate ASA on their way to a device. The echo packets will also pass through two of its interfaces as they return. This procedure can be used to perform a basic test of the interfaces, operation, and response time of the intermediate unit.

- Pinging to test questionable operation of a network device—A ping may be initiated from an ASA interface to a network device that is suspected of functioning incorrectly. If the interface is configured correctly and an echo is not received, there may be problems with the device.

- Pinging to test intermediate communications—A ping may be initiated from an ASA interface to a network device that is known to be functioning correctly and returning echo requests. If the echo is received, the correct operation of any intermediate devices and physical connectivity is confirmed.

Pinging From an ASA Interface

For basic testing of an interface, you can initiate a ping from an ASA interface to a network device that you know is functioning correctly and returning replies through the intermediate communications path. For basic testing, make sure you do the following:

- Verify receipt of the ping from the ASA interface by the "known good" device. If the ping is not received, a problem with the transmitting hardware or interface configuration may exist.

- If the ASA interface is configured correctly and it does not receive an echo reply from the "known good" device, problems with the interface hardware receiving function may exist. If a different interface with "known good" receiving capability can receive an echo after pinging the same "known good" device, the hardware receiving problem of the first interface is confirmed.

Pinging to an ASA Interface

When you try to ping to an ASA interface, verify that the pinging response (ICMP echo reply) is enabled for that interface by choosing **Tools > Ping**. When pinging is disabled, the ASA cannot be detected by other devices or software applications, and does not respond to the ASDM Ping tool.

Pinging Through the ASA Interface

To verify that other types of network traffic from "known good" sources are being passed through the ASA, choose **Monitoring > Interfaces > Interface Graphs** or an SNMP management station.

To enable internal hosts to ping external hosts, configure ICMP inspection. Choose **Configuration > Firewall > Service Policies**.

Troubleshooting the Ping Tool

When pings fail to receive an echo, it may be the result of a configuration or operational error in an ASA, and not necessarily because of no response from the IP address being pinged. Before using the Ping tool to ping from, to, or through an ASA interface, perform the following basic checks:

- Verify that interfaces are configured. Choose **Configuration > Device Setup > Interfaces**.

- Verify that devices in the intermediate communications path, such as switches or routers, are correctly delivering other types of network traffic.

- Make sure that traffic of other types from "known good" sources is being passed. Choose **Monitoring > Interfaces > Interface Graphs**.

Using the Ping Tool

To use the Ping tool, perform the following steps:

Step 1   In the main ASDM application window, choose **Tools > Ping**.

The Ping dialog box appears.

Step 2   Enter the destination IP address for the ICMP echo request packets in the IP Address field.

Ping also supports IPv6 addresses.

> ✎
>
> **Note**   If a hostname has been assigned in the Configuration > Firewall > Objects > Service Objects/Groups pane, you can use the hostname in place of the IP address.

Step 3   (Optional) Choose the ASA interface that transmits the echo request packets from the drop-down list. If it is not specified, the ASA checks the routing table to find the destination address and uses the required interface.

Step 4   Click **Ping** to send an ICMP echo request packet from the specified or default interface to the specified IP address and start the response timer.

The response appears in the Ping Output area. Three attempts are made to ping the IP address, and results display the following fields:

- The IP address of the device pinged or a device name, if available. The name of the device, if assigned, may be displayed, even if NO response is the result.

- When the ping is transmitted, a millisecond timer starts with a specified maximum, or timeout value. This timer is useful for testing the relative response times of different routes or activity levels.

- Example Ping output:

```
Sending 5, 100-byte ICMP Echos to out-pc, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

If the ping fails, the output is as follows:

```
Sending 5, 100-byte ICMP Echos to 10.132.80.101, timeout is 2 seconds:
?????
Success rate is 0 percent (0/5)
```

Step 5   To enter a new IP address, click **Clear Screen** to remove the previous response from the Ping output area.

Determining Packet Routing with Traceroute

The Traceroute tool helps you to determine the route that packets will take to their destination. The tool prints the result of each probe sent. Every line of output corresponds to a TTL value in increasing order. The following table lists the output symbols printed by this tool.

Output Symbol	Description
*	No response was received for the probe within the timeout period.
nn msec	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 unreachable.
!A	ICMP administratively prohibited.
?	Unknown ICMP error.

To use the Traceroute tool, perform the following steps:

Step 1 In the main ASDM application window, choose **Tools > Traceroute**.

The Traceroute dialog box appears.

Step 2 Enter hostname or IP address to which the route is traced. If the hostname is given, define it by choosing **Configuration > Firewall > Objects > Service Objects/Groups**, or configure a DNS server to enable this tool to resolve the hostname to an IP address.

Step 3 Enter the amount of time in seconds to wait for a response before the connection times out. The default is three seconds.

Step 4 Type the destination port used by the UDP probe messages. The default is 33434.

Step 5 Enter the number of probes to be sent at each TTL level. The default is three.

Step 6 Specify the minimum and maximum TTL values for the first probes. The minimum default is one, but it can be set to a higher value to suppress the display of known hops. The maximum default is 30. The traceroute terminates when the packet reaches the destination or when the maximum value is reached.

Step 7 Check the **Specify source interface or IP address** check box. Choose the source interface or IP address for the packet trace from the drop-down list. This IP address must be the IP address of one of the interfaces. In transparent mode, it must be the management IP address of the ASA.

Step 8 Check the **Reverse Resolve** check box to have the output display the names of hops encountered if name resolution is configured. Leave this check box unchecked to have the output display IP addresses.

Step 9 Check the **Use ICMP** check box to specify the use of ICMP probe packets instead of UDP probe packets.

Step 10 Click **Trace Route** to start the traceroute.

The Traceroute Output area displays detailed messages about the traceroute results.

Step 11 Click **Clear Output** to start a new traceroute.

Tracing Packets with Packet Tracer

The packet tracer tool provides packet tracing for packet sniffing and network fault isolation, as well as detailed information about the packets and how they are processed by the ASA. If a configuration command did not cause the packet to drop, the packet tracer tool can provide information about the cause in an easily readable format.

In addition, you can trace the lifespan of a packet through the ASA to see whether the packet is operating correctly with the packet tracer tool. This tool enables you to do the following:

- Debug all packet drops in a production network.
- Verify the configuration is working as intended.
- Show all rules applicable to a packet, along with the CLI commands that caused the rule addition.
- Show a time line of packet changes in a data path.
- Inject tracer packets into the data path.
- Search for an IPv4 or IPv6 address based on the user identity and the FQDN.

To use the packet tracer, perform the following steps:

Step 1   In the main ASDM application window, choose **Tools > Packet Tracer**.

The Cisco ASDM Packet Tracer dialog box appears.

Step 2   Choose the source interface for the packet trace from the drop-down list.

Step 3   Specify the protocol type for the packet trace. Available protocol types include ICMP, IP, TCP, and UDP.

Step 4   In the Source drop-down list, select one of the following options:

- IP Address
- User
- FQDN
- Security Tag
- Security Name

   Select the Security Tag or Security Name options when you want to trace packets sent by the ASA when integrated with the Cisco TrustSec solution. Security names are created on the Cisco ISE and provide user-friendly names for security groups.

   If a security policy is configured on the ASA with that security tags or security names, the ASA enforces the policy. (You can create security policies on the ASA that contain security tags or security names. To enforce policies based on security group names, the ASA needs the security group table to map security names to security tags.)

   See the general operations configuration guide for information about configuring the ASA to integrate with the Cisco TrustSec solution.

Step 5   Based on the option you selected from the Source drop-down list, enter the corresponding text for the item you want to trace; for example, enter the source IP address for the packet trace in the Source IP Address field.

Step 6   For TCP and UDP only, choose the source port for the packet trace from the drop-down list.

Step 7   In the Destination drop-down list, select one of the following options:

- IP Address
- FQDN

- Security Tag
- Security Name

Step 8   Based on the option you selected from the Destination drop-down list, enter the corresponding text for the item you want to trace; for example, enter the source IP address for the packet trace in the Destination IP Address field.

Step 9   For TCP and UDP only, choose the destination port for the packet trace from the drop-down list.

Step 10   For ICMP only, choose the type of packet trace from the Type drop-down list. Then enter the trace code and trace ID in the appropriate fields.

Step 11   For IP only, enter the protocol number in the Protocol field. Valid values range from 0 to 255.

Step 12   Click **Start** to trace the packet.

The Information Display Area shows detailed messages about the results of the packet trace.

> **Note**   To display a graphical representation of the packet trace, check the **Show animation** check box.

Step 13   Click **Clear** to start a new packet trace.

Monitoring Performance

To view ASA performance information in a graphical or tabular format, perform the following steps:

Step 1   In the ASDM main window, choose **Monitoring > Properties > Connection Graphs > Perfmon**.

Step 2   Select one or more entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. To remove an entry from the Selected Graphs list, click **Remove**. The available options are the following:

- AAA Perfmon—Displays the ASA AAA performance information.
- Inspection Perfmon—Displays the ASA inspection performance information.
- Web Perfmon—Displays the ASA web performance information, including URL access and URL server requests.
- Connections Perfmon—Displays the ASA connections performance information.
- Xlate Perfmon—Displays the ASA NAT performance information.

You can choose up to four types of statistics to show in one graph window. You can open multiple graph windows at the same time.

Step 3   To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title in the Graph Window Title field.

Step 4   Click **Show Graphs** to view performance statistics in a new or updated graph window.

Step 5   Click the **Table** tab to view the same performance statistics in a tabular format.

Step 6   From the View drop-down list on either tab, choose to display updates to information in the following time periods: Real-time, data every 10 sec; Last 10 minutes, data every 10 sec; Last 60 minutes, data every 1 min; Last 12 hours, data every 12 minutes; or Last 5 days, data every two hours.

Step 7   (Optional) Click **Export** to display the Export Graph Data dialog box. The selected performance statistics to export are already checked.

Step 8   (Optional) Click **Export** again to display the Save dialog box.

Step 9   (Optional) Click **Save** to save the performance statistics to a text file (.txt) on your local drive for future reference.

Step 10   (Optional) Click **Print** to display the Print Graph dialog box.

Step 11   (Optional) Choose the graph or table name from the drop-down list, then click **Print** to display the Print dialog box.

Step 12   (Optional) Click **OK** to print the selected performance statistics.

Monitoring System Resources

This section includes the following topics:

- Blocks, page 22-9
- CPU, page 22-10
- Memory, page 22-10

Blocks

To view the free and used memory blocks, perform the following steps:

Step 1   In the ASDM main window, choose **Monitoring > Properties > System Resources Graphs > Blocks**.

Step 2   Select one or more entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. To remove an entry from the Selected Graphs list, click **Remove**. The available options are the following:

- Blocks Used—Displays the ASA used memory blocks.
- Blocks Free—Displays the ASA free memory blocks.

You can choose up to four types of statistics to show in one graph window. You can open multiple graph windows at the same time.

Step 3   To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title in the Graph Window Title field.

Step 4   Click **Show Graphs** to view system resource statistics in a new or updated graph window.

Step 5   Click the **Table** tab to view the same performance statistics in a tabular format.

Step 6   From the View drop-down list on either tab, choose to display updates to information in the following time periods: Real-time, data every 10 sec; Last 10 minutes, data every 10 sec; Last 60 minutes, data every 1 min; Last 12 hours, data every 12 minutes; or Last 5 days, data every two hours.

Step 7   (Optional) Click **Export** to display the Export Graph Data dialog box. The selected memory block statistics to export are already checked.

Step 8   (Optional) Click **Export** again to display the Save dialog box.

Step 9 (Optional) Click **Save** to save the memory block statistics to a text file (.txt) on your local drive for future reference.

Step 10 (Optional) Click **Print** to display the Print Graph dialog box.

Step 11 (Optional) Choose the graph or table name from the drop-down list, then click **Print** to display the Print dialog box.

Step 12 (Optional) Click **OK** to print the selected memory block statistics.

CPU

To view the CPU utilization, perform the following steps:

Step 1 In the ASDM main window, choose **Monitoring > Properties > System Resources Graphs > CPU**.

Step 2 Select one or more entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. To remove an entry from the Selected Graphs list, click **Remove**.

You can choose up to four types of statistics to show in one graph window. You can open multiple graph windows at the same time.

Step 3 To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title in the Graph Window Title field.

Step 4 Click **Show Graphs** to view system resource statistics in a new or updated graph window.

Step 5 Click the **Table** tab to view the same performance statistics in a tabular format.

Step 6 From the View drop-down list on either tab, choose to display updates to information in the following time periods: Real-time, data every 10 sec; Last 10 minutes, data every 10 sec; Last 60 minutes, data every 1 min; Last 12 hours, data every 12 minutes; or Last 5 days, data every two hours.

Step 7 (Optional) Click **Export** to display the Export Graph Data dialog box. The selected CPU utilization statistics to export are already checked.

Step 8 (Optional) Click **Export** again to display the Save dialog box.

Step 9 (Optional) Click **Save** to save the CPU utilization statistics to a text file (.txt) on your local drive for future reference.

Step 10 (Optional) Click **Print** to display the Print Graph dialog box.

Step 11 (Optional) Choose the graph or table name from the drop-down list, then click **Print** to display the Print dialog box.

Step 12 (Optional) Click **OK** to print the selected CPU utilization statistics.

Memory

To view the memory utilization, perform the following steps:

Step 1 In the ASDM main window, choose **Monitoring > Properties > System Resources Graphs > Blocks**.

Step 2 Select one or more entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. To remove an entry from the Selected Graphs list, click **Remove**. The available options are the following:

- Free Memory—Displays the ASA free memory.
- Used Memory—Displays the ASA used memory.

You can choose up to four types of statistics to show in one graph window. You can open multiple graph windows at the same time.

Step 3 To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title in the Graph Window Title field.

Step 4 Click **Show Graphs** to view system resource statistics in a new or updated graph window.

Step 5 Click the **Table** tab to view the same performance statistics in a tabular format.

Step 6 From the View drop-down list on either tab, choose to display updates to information in the following time periods: Real-time, data every 10 sec; Last 10 minutes, data every 10 sec; Last 60 minutes, data every 1 min; Last 12 hours, data every 12 minutes; or Last 5 days, data every two hours.

Step 7 (Optional) Click **Export** to display the Export Graph Data dialog box. The selected memory utilization statistics to export are already checked.

Step 8 (Optional) Click **Export** again to display the Save dialog box.

Step 9 (Optional) Click **Save** to save the memory utilization statistics to a text file (.txt) on your local drive for future reference.

Step 10 (Optional) Click **Print** to display the Print Graph dialog box.

Step 11 (Optional) Choose the graph or table name from the drop-down list, then click **Print** to display the Print dialog box.

Step 12 (Optional) Click **OK** to print the selected memory utilization statistics.

Monitoring Connections

To view current connections in a tabular format, in the ASDM main window, choose **Monitoring > Properties > Connections**. Each connection is identified by the following parameters:

- Protocol
- Source:
  - Security ID
  - Security Name
  - IP address
  - Port
- Destination:
  - Security ID
  - Security Name
  - IP address
  - Port

- Idle time since the last packet was sent or received

- Amount of sent and received traffic on the connection

Monitoring Per-Process CPU Usage

You can monitor the processes that run on the CPU. You can obtain information about the percentage of CPU that is used by a certain process. CPU usage statistics are sorted in descending order to display the highest consumer at the top. Also included is information about the load on the CPU per process, at 5 seconds, 1 minute, and 5 minutes before the log time. This information is updated automatically every 5 seconds to provide real-time statistics. In ASDM, it is updated every 30 seconds.

To view CPU usage on a per-process basis, perform the following steps:

Step 1   In the ASDM main window, choose **Monitoring > Properties > Per-Process CPU Usage**.

Step 2   To pause the auto-refresh of the screen, click **Stop auto-refresh**.

Step 3   To save the information on the screen to a local text file, click **Save log to local file**.

The Save dialog box appears.

Step 4   Enter the name of the text file, then click **Save**.

To color code processes according to their CPU usage range, click **Configure CPU usage**.

The Color Settings dialog box appears.

Step 5   Choose one of the following range options: 49% and below, 50% to 79%, and 80% and above.

Step 6   Click the foreground or background cell to display the Pick a Color dialog box, and select the foreground and background colors for the given ranges.

Step 7   Click one of the following tabs to pick the color palette: **Swatches**, **HSB**, or **RGB**. When you are done, click **OK**.

Step 8   Click **OK** to view the color-coded entries.

Step 9   Click **Refresh** to refresh the data manually at any time.

P A R T   6

Advanced Network Protection

CHAPTER **23**

ASA and Cisco Cloud Web Security

Cisco Cloud Web Security provides web security and web filtering services through the Software-as-a-Service (SaaS) model. Enterprises with the ASA in their network can use Cloud Web Security services without having to install additional hardware.

When Cloud Web Security is enabled on the ASA, the ASA transparently redirects selected HTTP and HTTPS traffic to the Cloud Web Security proxy servers. The Cloud Web Security proxy servers then scan the content and allow, block, or send a warning about the traffic based on the policy configured in Cisco ScanCenter to enforce acceptable use and to protect users from malware.

The ASA can optionally authenticate and identify users with Identity Firewall (IDFW) and AAA rules. The ASA encrypts and includes the user credentials (including usernames and/or user groups) in the traffic it redirects to Cloud Web Security. The Cloud Web Security service then uses the user credentials to match the traffic to the policy. It also uses these credentials for user-based reporting. Without user authentication, the ASA can supply an (optional) default username and/or group, although usernames and groups are not required for the Cloud Web Security service to apply policy.

You can customize the traffic you want to send to Cloud Web Security when you create your service policy rules. You can also configure a "whitelist" so that a subset of web traffic that matches the service policy rule instead goes directly to the originally requested web server and is not scanned by Cloud Web Security.

You can configure a primary and a backup Cloud Web Security proxy server, each of which the ASA polls regularly to check for availability.

Note This feature is also called "ScanSafe," so the ScanSafe name appears in some commands.

This chapter includes the following sections:

Information About Cisco Cloud Web Security

This section includes the following topics:

Redirection of Web Traffic to Cloud Web Security

When an end user sends an HTTP or HTTPS request, the ASA receives it and optionally retrieves the user and/or group information. If the traffic matches an ASA service policy rule for Cloud Web Security, then the ASA redirects the request to the Cloud Web Security proxy servers. The ASA acts as an intermediary between the end user and the Cloud Web Security proxy server by redirecting the connection to the proxy server. The ASA changes the destination IP address and port in the client requests and adds Cloud Web Security-specific HTTP headers and then sends the modified request to the Cloud Web Security proxy server. The Cloud Web Security HTTP headers include various kinds of information, including the username and user group (if available).

User Authentication and Cloud Web Security

User identity can be used to apply policy in Cloud Web Security. User identity is also useful for Cloud Web Security reporting. User identity is not required to use Cloud Web Security. There are other methods to identify traffic for Cloud Web Security policy.

The ASA supports the following methods of determining the identity of a user, or of providing a default identity:

- AAA rules—When the ASA performs user authentication using a AAA rule, the username is retrieved from the AAA server or local database. Identity from AAA rules does not include group information. If configured, the default group is used. For information about configuring AAA rules, see the legacy feature guide.

- IDFW—When the ASA uses IDFW with the Active Directory (AD), the username and group is retrieved from the AD agent when you activate a user and/or group by using an ACL in a feature such as an access rule or in your service policy, or by configuring the user identity monitor to download user identity information directly.

  For information about configuring IDFW, see the general operations configuration guide.

- Default username and group—Without user authentication, the ASA uses an optional default username and/or group for all users that match a service policy rule for Cloud Web Security.

Authentication Keys

Each ASA must use an authentication key that you obtain from Cloud Web Security. The authentication key lets Cloud Web Security identify the company associated with web requests and ensures that the ASA is associated with valid customer.

You can use one of two types of authentication keys for your ASA: the company key or the group key.

- Company Authentication Key, page 23-3
- Group Authentication Key, page 23-3

Company Authentication Key

A Company authentication key can be used on multiple ASAs within the same company. This key simply enables the Cloud Web Security service for your ASAs. The administrator generates this key in ScanCenter (https://scancenter.scansafe.com/portal/admin/login.jsp); you have the opportunity to e-mail the key for later use. You cannot look up this key later in ScanCenter; only the last 4 digits are shown in ScanCenter. For more information, see the Cloud Web Security documentation: http://www.cisco.com/en/US/products/ps11720/products_installation_and_configuration_guides_list.html.

Group Authentication Key

A Group authentication key is a special key unique to each ASA that performs two functions:

- Enables the Cloud Web Security service for one ASA.
- Identifies all traffic from the ASA so you can create ScanCenter policy per ASA.

For information about using the Group authentication key for policy, see ScanCenter Policy, page 23-4).

The administrator generates this key in ScanCenter (https://scancenter.scansafe.com/portal/admin/login.jsp); you have the opportunity to e-mail the key for later use. You cannot look up this key later in ScanCenter; only the last 4 digits are shown in ScanCenter. For more information, see the Cloud Web Security documentation: http://www.cisco.com/en/US/products/ps11720/products_installation_and_configuration_guides_list.html.

ScanCenter Policy

In ScanCenter, traffic is matched against policy rules in order until a rule is matched. Cloud Web Security then applies the configured action for the rule. User traffic can match a policy rule in ScanCenter based on group association: a *directory group* or a *custom group*.

Directory Groups

Directory groups define the group to which traffic belongs. The group, if present, is included in the HTTP header of the client request. The ASA includes the group in the HTTP header when you configure IDFW. If you do not use IDFW, you can configure a default group for traffic matching an ASA rule for Cloud Web Security inspection.

When you configure a directory group, you must enter the group name exactly.

- IDFW group names are sent in the following format:

  domain-name\group-name

  When the ASA learns the IDFW group name, the format on the ASA is *domain-name\\group-name*. However, the ASA modifies the name to use only one backslash (\) to conform to typical ScanCenter notation.

- The default group name is sent in the following format:

  [*domain*]*group-name*

  On the ASA, you need to configure the optional domain name to be followed by 2 backslashes (\\); however, the ASA modifies the name to use only one backslash (\) to conform to typical ScanCenter notation. For example, if you specify "Cisco\\Boulder1," the ASA modifies the group name to be "Cisco\Boulder1" with only one backslash (\) when sending the group name to Cloud Web Security.

Custom Groups

Custom groups are defined using one or more of the following criteria:

- ScanCenter Group authentication key—You can generate a Group authentication key for a custom group. Then, if you identify this group key when you configure the ASA, all traffic from the ASA is tagged with the Group key.

- Source IP address—You can identify source IP addresses in the custom group. Note that the ASA service policy is based on source IP address, so you might want to configure any IP address-based policy on the ASA instead.

- Username—You can identify usernames in the custom group.
  - IDFW usernames are sent in the following format:

    domain-name\username

  - AAA usernames, when using RADIUS or TACACS+, are sent in the following format:

    LOCAL*username*

  - AAA usernames, when using LDAP, are sent in the following format:

    domain-name\username

 – For the default username, it is sent in the following format:

 [*domain-name*]*username*

 For example, if you configure the default username to be "Guest," then the ASA sends "Guest." If you configure the default username to be "Cisco\Guest," then the ASA sends "Cisco\Guest."

How Groups and the Authentication Key Interoperate

Unless you need the per-ASA policy that a custom group+group key provides, you will likely use a company key. Note that not all custom groups are associated with a group key. Non-keyed custom groups can be used to identify IP addresses or usernames, and can be used in your policy along with rules that use directory groups.

Even if you do want per-ASA policy and are using a group key, you can also use the matching capability provided by directory groups and non-keyed custom groups. In this case, you might want an ASA-based policy, with some exceptions based on group membership, IP address, or username. For example, if you want to exempt users in the America\Management group across all ASAs:

1. Add a directory group for America\Management.

2. Add an exempt rule for this group.

3. Add rules for each custom group+group key after the exempt rule to apply policy per-ASA.

4. Traffic from users in America\Management will match the exempt rule, while all other traffic will match the rule for the ASA from which it originated.

Many combinations of keys, groups, and policy rules are possible.

Cloud Web Security Actions

After applying the configured policies, Cloud Web Security either blocks, allows, or sends a warning about the user request:

 • Allows—When Cloud Web Security allows the client request, it contacts the originally requested server and retrieves the data. It forwards the server response to the ASA, which then forwards it to the user.

 • Blocks—When Cloud Web Security blocks the client request, it notifies the user that access has been blocked. It sends an HTTP 302 "Moved Temporarily" response that redirects the client application to a web page hosted by the Cloud Web Security proxy server showing the blocked error message. The ASA forwards the 302 response to the client.

 • Warns—When the Cloud Web Security proxy server determines that a site may be in breach of the acceptable use policy, it displays a warning page about the site. You can choose to heed the warning and drop the request to connect, or you can click through the warning and proceed to the requested site.

You can also choose how the ASA handles web traffic when it cannot reach either the primary or backup Cloud Web Security proxy server. It can block or allow all web traffic. By default, it blocks web traffic.

Bypassing Scanning with Whitelists

If you use AAA rules or IDFW, you can configure the ASA so that web traffic from specific users or groups that otherwise match the service policy rule is not redirected to the Cloud Web Security proxy server for scanning. When you bypass Cloud Web Security scanning, the ASA retrieves the content

directly from the originally requested web server without contacting the proxy server. When it receives the response from the web server, it sends the data to the client. This process is called "whitelisting" traffic.

Although you can achieve the same results of exempting traffic based on user or group when you configure the class of traffic using ACLs to send to Cloud Web Security, you might find it more straightforward to use a whitelist instead. Note that the whitelist feature is only based on user and group, not on IP address.

IPv4 and IPv6 Support

Cloud Web Security currently supports only IPv4 addresses. If you use IPv6 internally, NAT 64 must be performed for any IPv6 flows that need to be sent to Cloud Web Security.

The following table shows the class map traffic that is supported by Cloud Web Security redirection:

Class Map Traffic	Cloud Web Security Inspection
From IPv4 to IPv4	Supported
From IPv6 to IPv4 (using NAT64)	Supported
From IPv4 to IPv6	Not Supported
From IPv6 to IPv6	Not Supported

Failover from Primary to Backup Proxy Server

When you subscribe to the Cisco Cloud Web Security service, you are assigned a primary Cloud Web Security proxy server and backup proxy server.

If any client is unable to reach the primary server, then the ASA starts polling the tower to determine availability. (If there is no client activity, the ASA polls every 15 miniutes.) If the proxy server is unavailable after a configured number of retries (the default is 5; this setting is configurable), the server is declared unreachable, and the backup proxy server becomes active.

If a client or the ASA can reach the server at least twice consecutively before the retry count is reached, the polling stops and the tower is determined to be reachable.

After a failover to the backup server, the ASA continues to poll the primary server. If the primary server becomes reachable, then the ASA returns to using the primary server.

Licensing Requirements for Cisco Cloud Web Security

Model	License Requirement
ASAv	Standard or Premium License.
All other models	Strong Encryption (3DES/AES) License to encrypt traffic between the security appliance and the Cloud Web Security server.

On the Cloud Web Security side, you must purchase a Cisco Cloud Web Security license and identify the number of users that the ASA handles. Then log into ScanCenter, and generate your authentication keys.

Prerequisites for Cloud Web Security

(Optional) User Authentication Prerequisites

To send user identity information to Cloud Web Security, configure one of the following on the ASA:

- AAA rules (username only)—See the legacy feature guide.
- IDFW (username and group)—See the general operations configuration guide.

(Optional) Fully Qualified Domain Name Prerequisites

If you use FQDNs in ACLs for your service policy rule, or for the Cloud Web Security server, you must configure a DNS server for the ASA according to the general operations configuration guide.

Guidelines and Limitations

Context Mode Guidelines

Supported in single and multiple context modes.

In multiple context mode, the server configuration is allowed only in the system, and the service policy rule configuration is allowed only in the security contexts.

Each context can have its own authentication key, if desired.

Firewall Mode Guidelines

Supported in routed firewall mode only. Does not support transparent firewall mode.

IPv6 Guidelines

Does not support IPv6. See .

Additional Guidelines

- Cloud Web Security is not supported with ASA clustering.
- Clientless SSL VPN is not supported with Cloud Web Security; be sure to exempt any clientless SSL VPN traffic from the ASA service policy for Cloud Web Security.
- When an interface to the Cloud Web Security proxy servers goes down, output from the **show scansafe server** command shows both servers up for approximately 15-25 minutes. This condition may occur because the polling mechanism is based on the active connection, and because that interface is down, it shows zero connection, and it takes the longest poll time approach.
- Cloud Web Security is not supported with the ASA CX module. If you configure both the ASA CX action and Cloud Web Security inspection for the same traffic, the ASA only performs the ASA CX action.
- Cloud Web Security inspection is compatibile with HTTP inspection for the same traffic. HTTP inspection is enabled by default as part of the default global policy.

- Cloud Web Security is not supported with extended PAT or any application that can potentially use the same source port and IP address for separate connections. For example, if two different connections (targeted to separate servers) use extended PAT, the ASA might reuse the same source IP and source port for both connection translations because they are differentiated by the separate destinations. When the ASA redirects these connections to the Cloud Web Security server, it replaces the destination with the Cloud Web Security server IP address and port (8080 by default). As a result, both connections now appear to belong to the same flow (same source IP/port and destination IP/port), and return traffic cannot be untranslated properly.

- The Default Inspection Traffic traffic class does not include the default ports for the Cloud Web Security inspection (80 and 443).

Default Settings

By default, Cisco Cloud Web Security is not enabled.

Configuring Cisco Cloud Web Security

- Configuring Communication with the Cloud Web Security Proxy Server, page 23-8
- (Multiple Context Mode) Allowing Cloud Web Security Per Security Context, page 23-10
- Configuring a Service Policy to Send Traffic to Cloud Web Security, page 23-10
- (Optional) Configuring Whitelisted Traffic, page 23-23
- Configuring the Cloud Web Security Policy, page 23-26

Configuring Communication with the Cloud Web Security Proxy Server

Guidelines

The public key is embedded in the ASA software, so there is no need for you to configure it.

Detailed Steps

Step 1   Choose **Configuration > Device Management > Cloud Web Security**.

Step 2   In the Primary Server area, enter the following:

- **•** IP Address/Domain Name—Enter the IPv4 address or FQDN of the primary server.

- **•** HTTP Port—Enter the HTTP port of the primary server (port to which traffic must be redirected). By default the port is 8080; do not change this value unless directed to do so.

Step 3   In the Backup Server area, enter the following:

- **•** IP Address/Domain Name—Enter the IPv4 address or FQDN of the backup server.

- **•** HTTP Port—Enter the HTTP port of the backup server (port to which traffic must be redirected). By default the port is 8080. Valid values are from 1 to 65535.

Step 4   In the Other area, enter the following:

- **•** Retry Counter—Enter the value for the number of consecutive polling failures to the Cloud Web Security proxy server before determining the server is unreachable. Polls are performed every 30 seconds. Valid values are from 2 to 100, and the default is 5.

- **•** License Key—Configure the authentication key that the ASA sends to the Cloud Web Security proxy servers to indicate from which organization the request comes. The authentication key is a 16-byte hexidecimal number. See Authentication Keys, page 23-3.

- **•** Confirm License Key—Confirm the authentication key.

Step 5   Click **Apply**.

(Multiple Context Mode) Allowing Cloud Web Security Per Security Context

In multiple context mode, you must allow Cloud Web Security per context. For more information, see the general operations configuration guide.

✎ **Note**   You must configure a route pointing to the Scansafe towers in both; the admin context and the specific context. This ensures that the Scansafe tower does not become unreachable in the Active/Active failover scenario.

Configuring a Service Policy to Send Traffic to Cloud Web Security

Your service policy consists of multiple service policy rules, applied globally, or applied to each interface. Each service policy rule can either send traffic to Cloud Web Security (Match) or exempt traffic from Cloud Web Security (Do Not Match). Create rules for traffic destined for the Internet. The order of these rules is important. When the ASA decides whether to forward or exempt a packet, the ASA tests the packet with each rule in the order in which the rules are listed. After a match is found, no more rules are checked. For example, if you create a rule at the beginning of a policy that explicitly Matches all traffic, no further statements are ever checked. You can reorder the rules as needed after you add them.

See Chapter 1, "Service Policy," for more information about service policy rules.

Prerequisites

(Optional) If you need to use a whitelist to exempt some traffic from being sent to Cloud Web Security, first create the whitelist according to the (Optional) Configuring Whitelisted Traffic, page 23-23 so you can refer to the whitelist in your service policy rule.

Detailed Steps

Step 1   Choose **Configuration > Firewall > Service Policy Rules**, and click **Add > Service Policy Rule** to add a service policy rule.

Step 2   On the Service Policy dialog box, you can configure Cloud Web Security as part of a new service policy, or you can edit an existing service policy. Click **Next**.
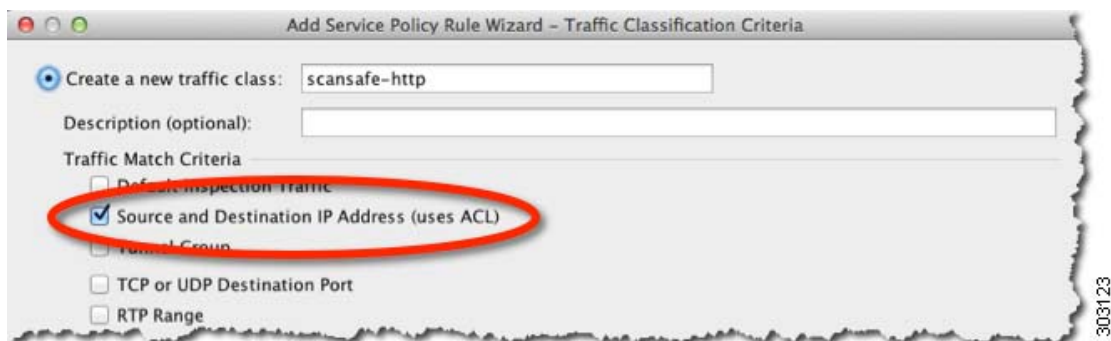
Step 3   On the Traffic Classification Criteria dialog box, name the traffic class (or accept the default name), keep the **Create a new traffic class** option selected, and click **Source and Destination IP address (Uses ACL)**, then click **Next**.

When you create a new traffic class of this type, you can only specify one access control entry (ACE) initially. After you finish adding the rule, you can add additional ACEs by adding a new rule to the same interface or global policy, and then specifying **Add rule to existing traffic class** on the Traffic Classification dialog box.

The Traffic Match - Source and Destination dialog box appears.
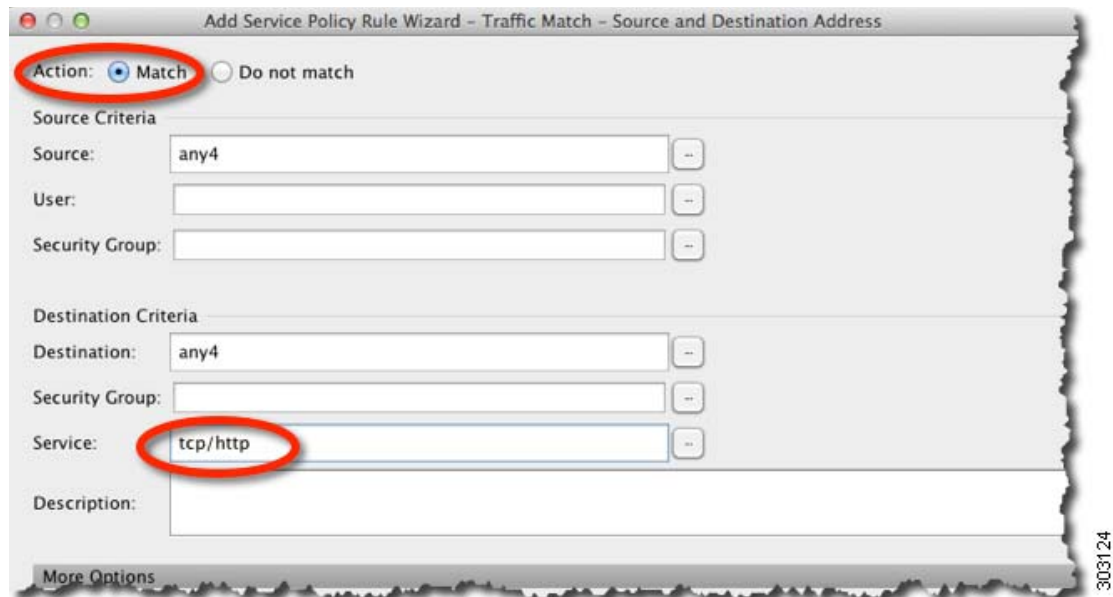
a. Click **Match** or **Do Not Match**.

   Match specifies that traffic matching the source and destination is sent to Cloud Web Security. **Do Not Match** exempts matching traffic from Cloud Web Security. You can later add additional rules to match or not match other traffic.

   When creating your rules, consider how you can match appropriate traffic that is destined for the Internet, but not match traffic that is destined for other internal networks. For example, to prevent inside traffic from being sent to Cloud Web Security when the destination is an internal server on the DMZ, be sure to add a deny ACE to the ACL that exempts traffic to the DMZ.

b. In the Source Criteria area, enter or browse for a Source IP address or network object, an optional IDFW Username or group, and an optional TrustSec Security Group.

c. In the Destination Criteria area, enter or browse for a Destination IP address or network object, and an optional TrustSec Security Group.

   FQDN network objects might be useful in matching or exempting traffic to specific servers.

d. In the Service field, enter **http** or **https**, and click **Next**.

   ✎

   Note   Cloud Web Security only operates on HTTP and HTTPS traffic. Each type of traffic is treated separately by the ASA. Therefore, you need to create HTTP-only rules and HTTPS-only rules.

   The Rule Actions dialog box appears.

Step 4   On the Protocol Inspection tab, check the **Cloud Web Security** check box.

Step 5   Click **Configure** to set the traffic action (fail open or fail close) and add the inspection policy map.

The inspection policy map configures essential parameters for the rule and also optionally identifies the whitelist. An inspection policy map is required for each class of traffic that you want to send to Cloud Web Security. You can also pre-configure inspection policy maps from the Configuration > Firewall > Objects > Inspect Maps > Cloud Web Security pane.

The Select Cloud Web Security Inspect Map dialog box appears.

a.   For the Cloud Web Security Traffic Action, choose one:

   –   F**ail Close**—Drops all traffic if the Cloud Web Security servers are unavailable.

   –   **Fail Open**—Allows traffic to pass through the ASA if the Cloud Web Security servers are unavailable.

b.   Choose an existing inspection policy map, or add one using the **Add** button.

c.   Click **Add** to add a new inspection policy map.

The Add Cloud Web Security Inspect Map dialog box appears.

d. In the Name field, specify a name for the inspection policy map, up to 40 characters in length.

e. (Optional) Enter a description.

f. (Optional) On the Parameters tab, specify a Default User and/or a Default Group. If the ASA cannot determine the identity of the user coming into the ASA, then the default user and/or group is applied.

g. For the Protocol, click **HTTP** or **HTTPS**, to match the service you set in Step 3d. Cloud Web Security treats each type of traffic separately.

h. (Optional) To identify a whitelist, click the **Inspections** tab.
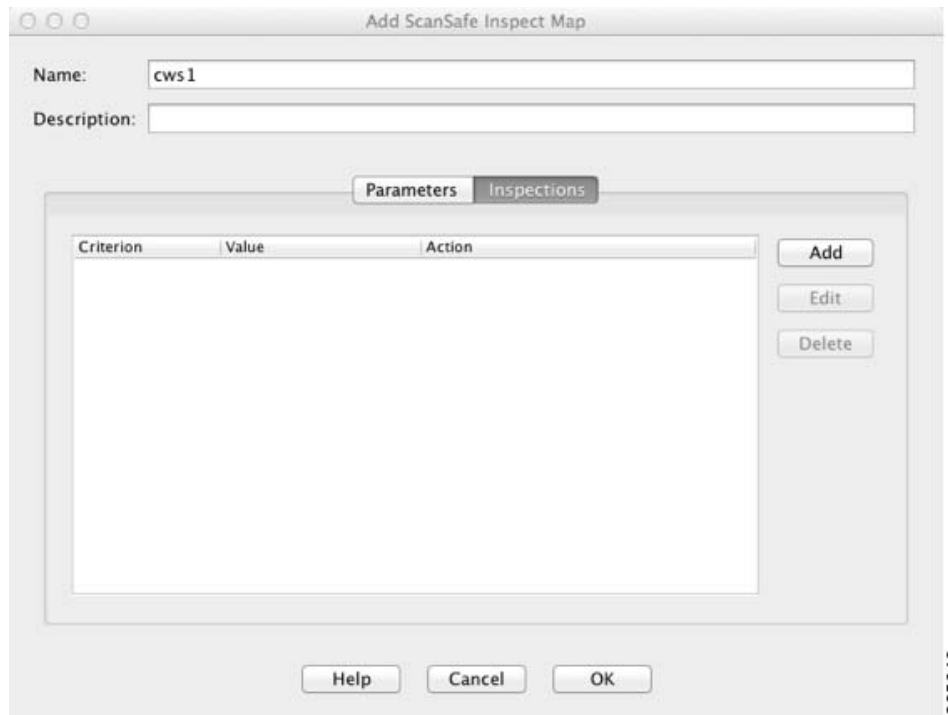
- **–** Click **Add** to choose the inspection class map you created in the (Optional) Configuring Whitelisted Traffic, page 23-23.

  The Add Cloud Web Security Match Criterion dialog box appears.

- **–** From the Cloud Web Security Traffic Class drop-down menu, choose an inspection class map.

  To add or edit a class map, click **Manage**.

- **–** For the Action, click **Whitelist**.

- **–** Click **OK** to add the whitelist to the policy map.

- **–** Click **OK**.

Step 6   Click **Finish**. The rule is added to the Service Policy Rules table.

Step 7   To add additional sub-rules (ACEs) for this traffic class, to match or exempt additional traffic:

    a.   Choose **Configuration > Firewall > Service Policy Rules**, and click **Add > Service Policy Rule**.

    b.   Choose the same service policy as Step 2. Click **Next**.
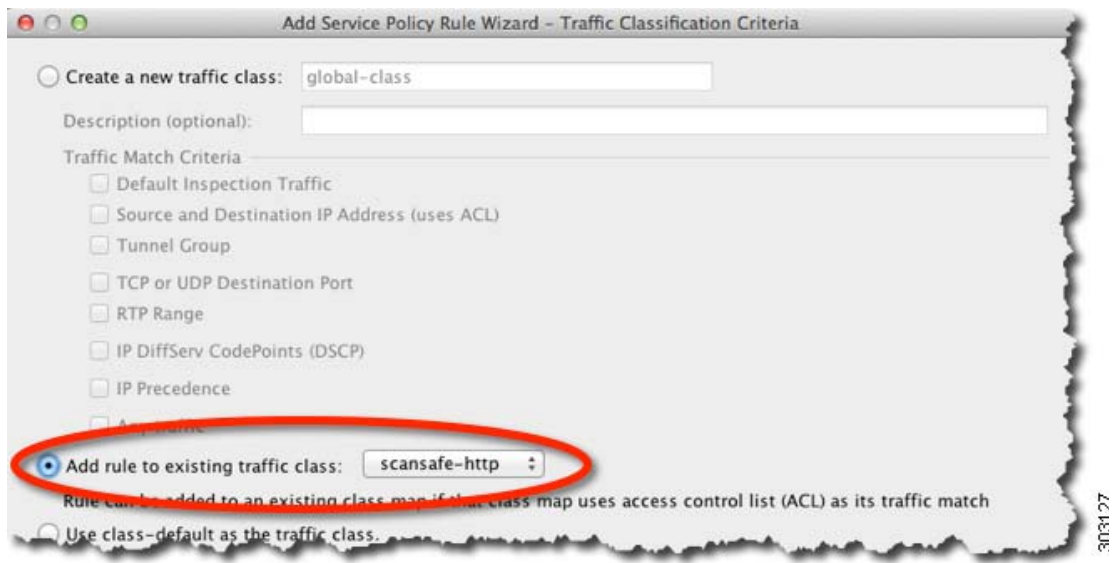
c. On the Traffic Classification Criteria dialog box, choose **Add Rule to Existing Traffic Class**, and choose the name you created in Step 3. Click **Next**.

d. In the Traffic Match - Source and Destination dialog box, choose **Match** to add inspect additional traffic, or **Do Not Match** to exempt traffic from Cloud Web Security inspections. Be sure to set the service to match the previous rules in this class (HTTP or HTTPS); you cannot mix HTTP and HTTPS in the same traffic class for Cloud Web Security. Click **Next**.

e. On the Rule Actions dialog box, do not make any changes; click **Finish**. For this traffic class, you can have only one set of rule actions even if you add multiple ACEs, so the previously-specified actions are inherited.

Step 8   Repeat this entire procedure to create an additional traffic class, for example for HTTPS traffic. You can create as many rules and sub-rules as needed.

Step 9   Arrange the order of Cloud Web Security rules and sub-rules on the Service Policy Rules pane. See Managing the Order of Service Policy Rules, page 1-15 for information about changing the order of ACEs.
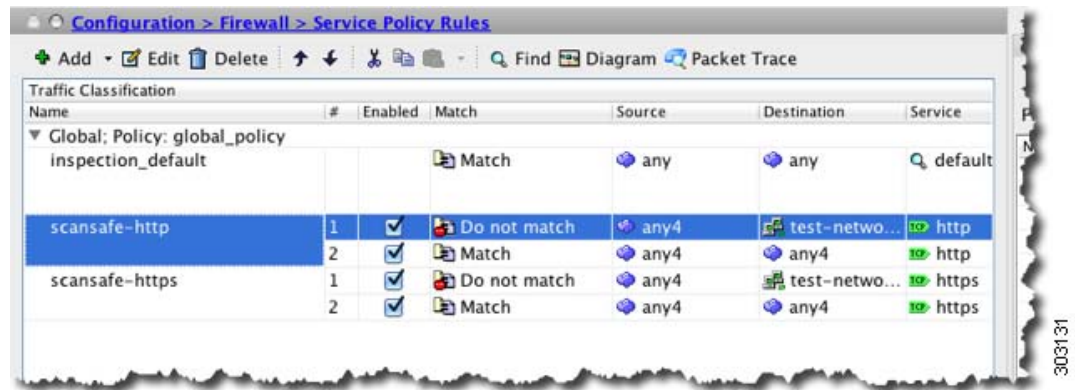
Step 10 Click **Apply**.

Examples

The following example exempts all IPv4 HTTP and HTTPS traffic going to the 10.6.6.0/24 (test_network), and sends all other HTTPS and HTTPS traffic to Cloud Web Security, and applies this service policy rule to all interfaces as part of the existing global policy. If the Cloud Web Security server is unreachable, the ASA drops all matching traffic (fail close). If a user is not have user identity information, the default user Boulder and group Cisco is used.

Step 1 Choose **Configuration > Firewall > Service Policy Rules**, and click **Add > Service Policy Rule**. Add this rule to the default global_policy:
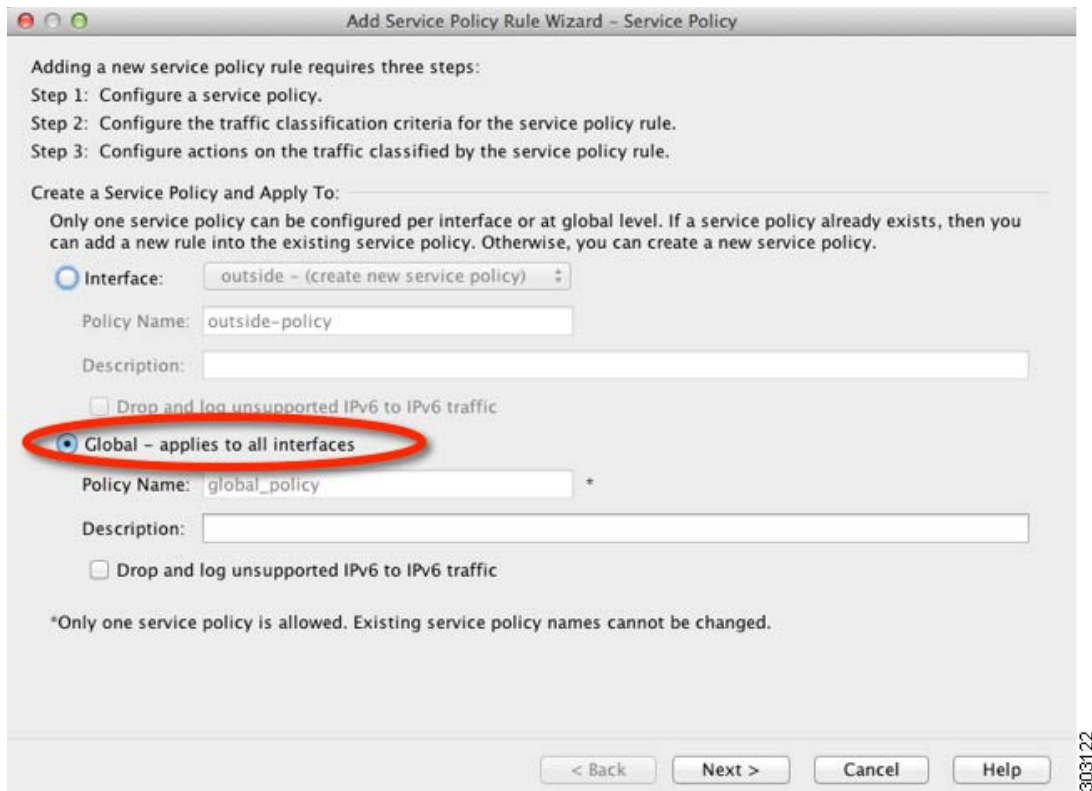
Step 2 Add a new traffic class called "scansafe-http," and specify an ACL for traffic matching:
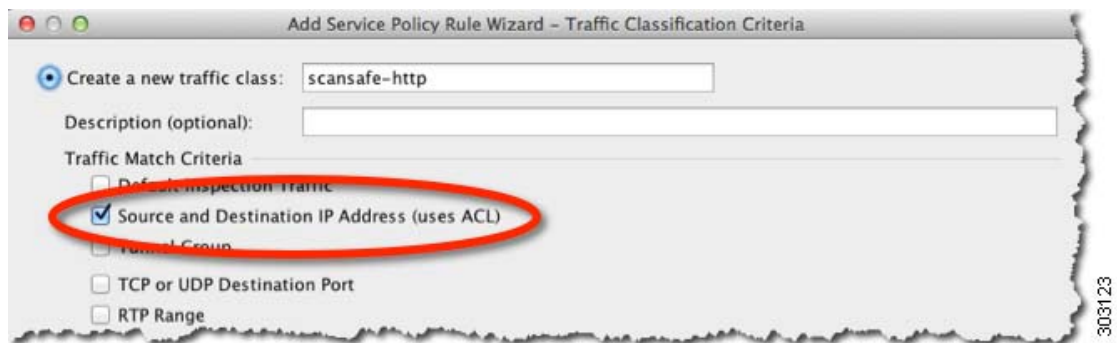
Step 3 Choose **Match**, and specify **any4** for the Source and Destination. Specify **tcp/http** for the Service.

Step 4   Check **Cloud Web Security** and click **Configure**.

Step 5   Accept the default Fail Close action, and click **Add**.

Step 6   Name the inspection policy map "http-map," set the Default User to Boulder and the default group to Cisco. Choose **HTTP**.

Step 7   Click **OK**, **OK**, and then **Finish**. The rule is added to the Service Policy Rules table.

Step 8   Choose **Configuration > Firewall > Service Policy Rules**, and click **Add > Service Policy Rule**. Add the new rule to the default global_policy:
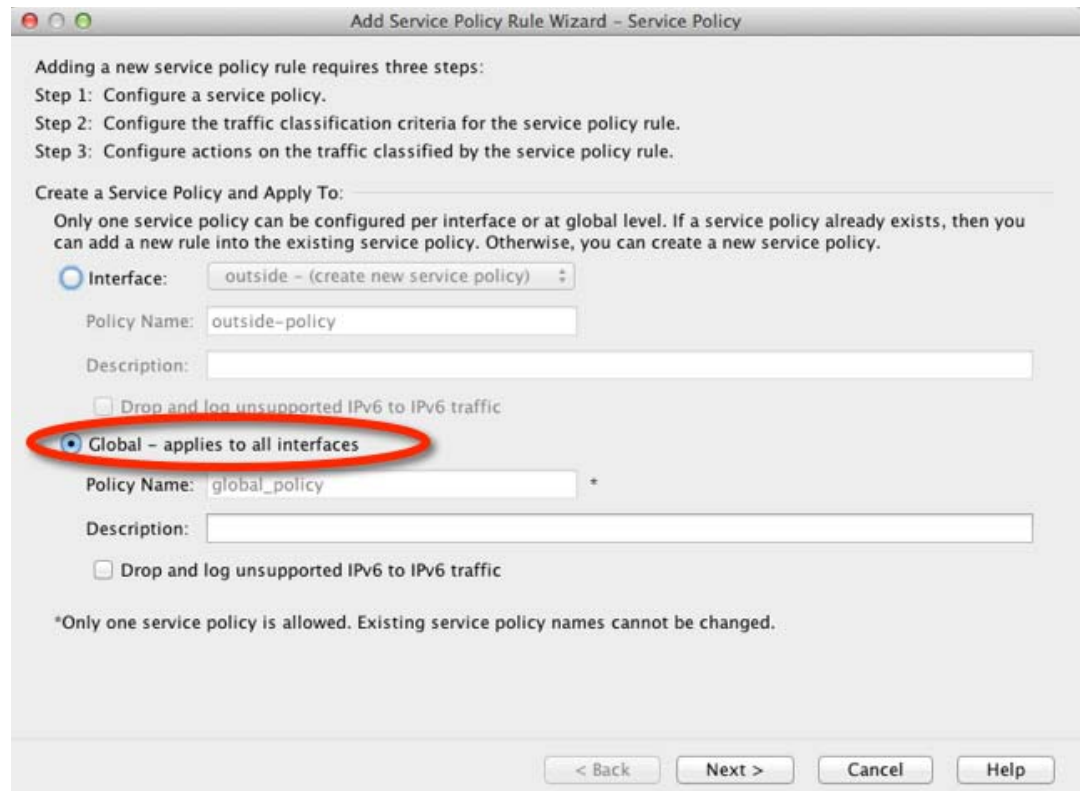
Step 9   Click **Add rule to existing traffic class**, and choose **scansafe-http**.
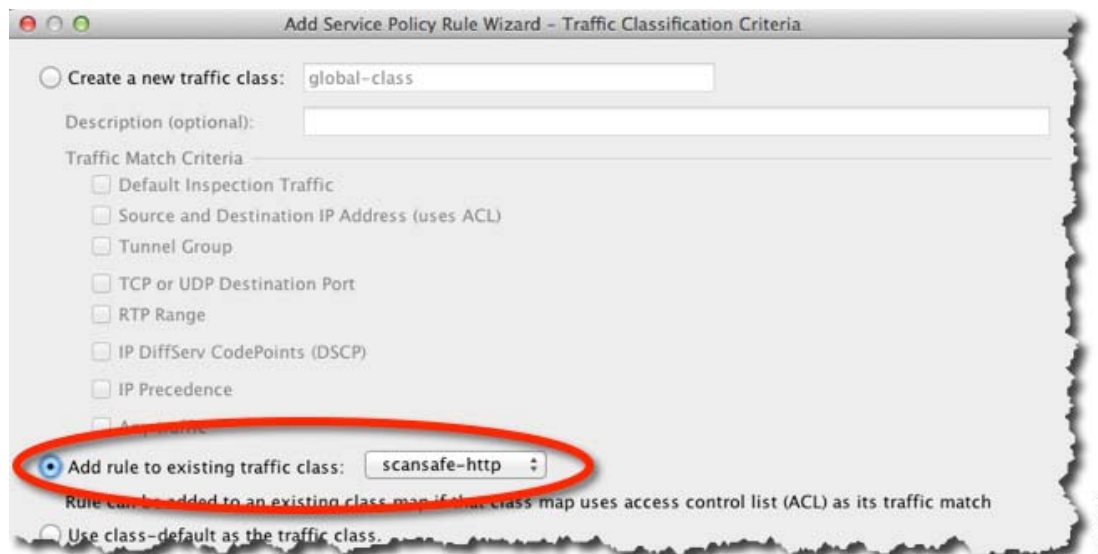
Step 10   Choose **Do not match**, set **any4** as the Source, and **10.6.6.0/24** as the Destination. Set the Service to **tcp/http**.

Step 11 Click **Finish**.
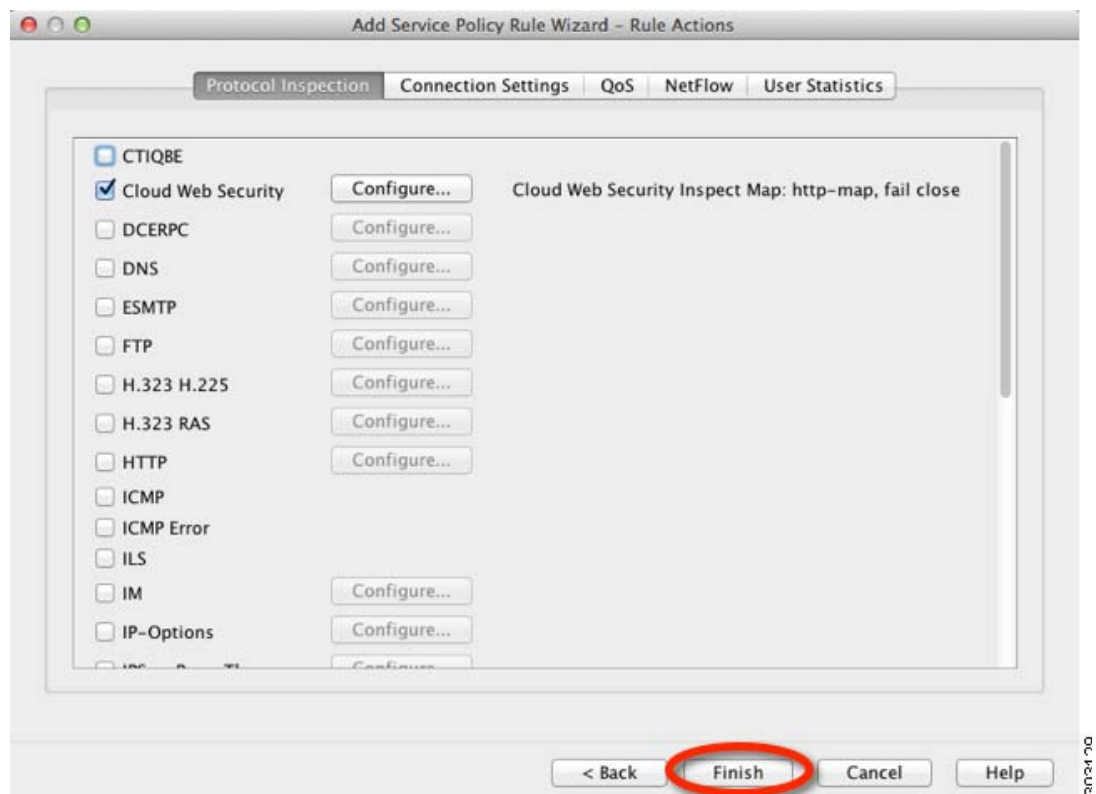
Step 12 Reorder the rules so the Do not match rule is above the Match rule.
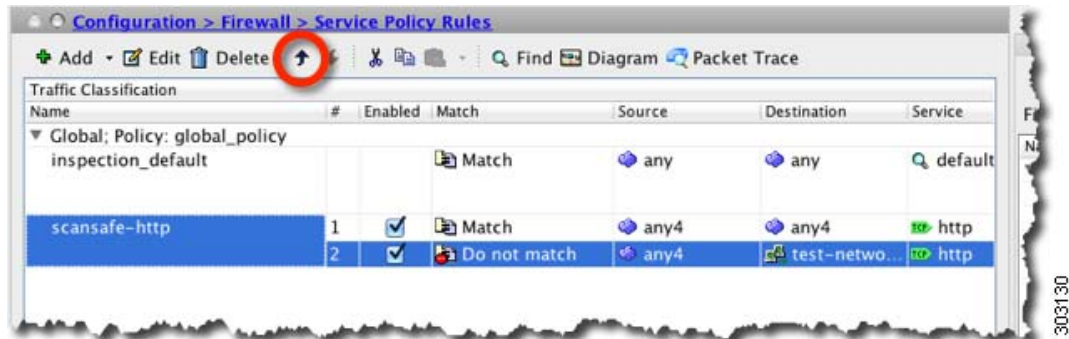
User traffic is compared to these rules in order; if this Match rule is first in the list, then all traffic, including traffic to test_network, will match only that rule and the Do not match rule will never be hit. If you move the Do not match rule above the Match rule, then traffic to test_network will match the Do not match rule, and all other traffic will match the Match rule.

Step 13   Repeat the above steps with the following changes: add a new traffic class called "scansafe-https," and choose **HTTPS** for the inspection policy map.
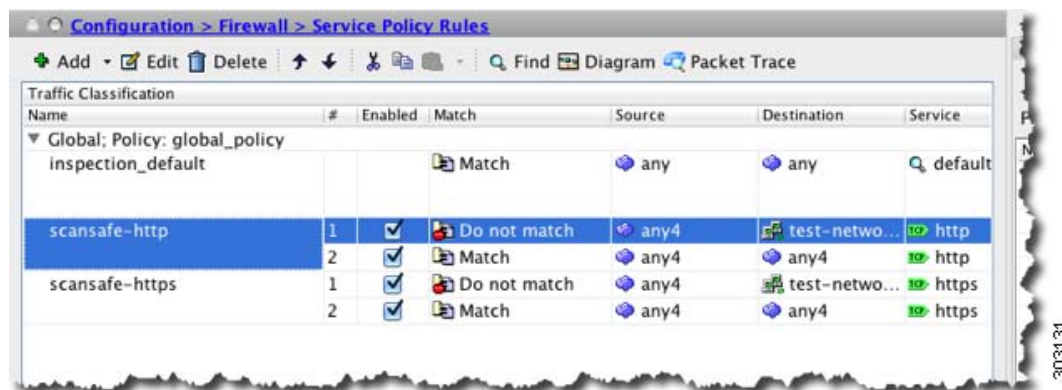
Step 14   Click **Apply**.

(Optional) Configuring Whitelisted Traffic

If you use user authentication, you can exempt some traffic from being filtered by Cloud Web Security based on the username and/or groupname. When you configure your Cloud Web Security service policy rule, you can reference the whitelisting inspection class map. Both IDFW and AAA user credentials can be used with this feature.

Although you can achieve the same results of exempting traffic based on user or group when you configure the service policy rule, you might find it more straightforward to use a whitelist instead. Note that the whitelist feature is only based on user and group, not on IP address.

Detailed Steps

Step 1   Choose **Configuration > Firewall > Objects > Class Maps > Cloud Web Security**.

Step 2   Click **Add** to create a new class map.

The Add Cloud Web Security Traffic Class Map screen appears.
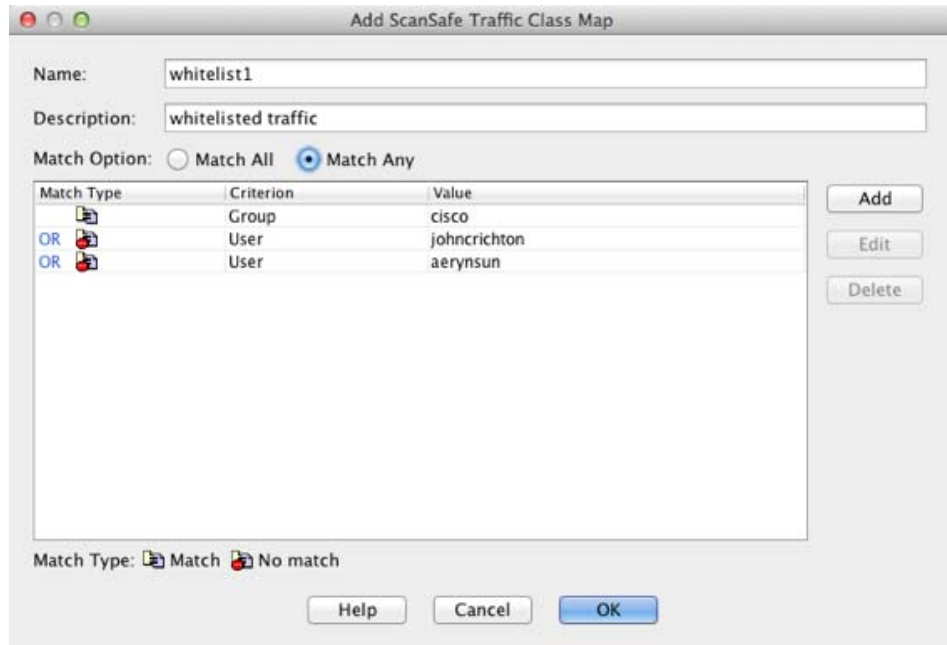
Step 3   In the Name field, enter the name of the new class map (40 characters or less).

Step 4   In the Description field, provide a description for the class map (200 characters or less).

Step 5   Choose the Match Option for the criteria you define when you click ADD:

- Match All—Specifies that traffic must match all criteria to match the class map.

- Match Any—Specifies that the traffic matches the class map if it matches at least one of the criteria.

Step 6   Click **Add**.

The Add Cloud Web Security Match Criterion Window appears.

Step 7   Choose the Match Type:

- Match—Specifies the user and/or group that you want to whitelist.

- No Match—Specifies the user and/or group that you do *not* want to whitelist; for example, if you whitelist the group "cisco," but you want to scan traffic from users "johncrichton" and "aerynsun," you can specify No Match for those users.

Step 8   Choose the Match Criterion:

- User—Specifies the user.

- Group—Specifies the group.

- User and Group—Specifies a user and group.

Step 9   Click **OK**.

Step 10   Continue to add match criteria as desired.

Step 11   Click **OK** to add the class map.

Step 12   Click **Apply**.

Step 13   Use the whitelist in the Cloud Web Security policy according to the Configuring a Service Policy to Send Traffic to Cloud Web Security, page 23-10.

(Optional) Configuring the User Identity Monitor

When you use IDFW, the ASA only downloads user identity information from the AD server for users and groups included in active ACLs; the ACL must be used in a feature such as an access rule, AAA rule, service policy rule, or other feature to be considered active. Because Cloud Web Security can base its policy on user identity, you may need to download groups that are not part of an active ACL to get full IDFW coverage for all your users. For example, although you can configure your Cloud Web Security service policy rule to use an ACL with users and groups, thus activating any relevant groups, it is not required; you could use an ACL based entirely on IP addresses.The user identity monitor feature lets you download group information directly from the AD agent.

Restrictions

The ASA can only monitor a maximum of 512 groups, including those configured for the user identity monitor and those monitored through active ACLs.

Detailed Steps

Step 1   Choose **Configuration > Firewall > Identity Options**, and scroll to the Cloud Web Security Configuration section.

Step 2   Click **Add**.

The Add Monitor User dialog box appears.

Step 3   To add a domain, click **Manage**, and then click **Add**. You can only monitor groups for domains you have pre-defined on the ASA.

The Configure Identity Domains dialog box appears. For detailed information about adding domains, see the general operations configuration guide.

Step 4   When you are finished adding domains, click **OK**.

Step 5   You can either type in a group name, or you can search for groups on the AD agent per domain.

- To type in a group name directly, enter the name in the bottom field in the following format, and click **OK**:

    domain-name*group*

- To search for a group on the AD agent:

a.   Choose the domain from the Domain drop-down list.

b.   In the Find field, enter a text string to match group names, and click **Find**.

The ASA downloads names from the AD agent for the specified domain.

c.   Double-click the name you want to monitor; it is added to the bottom field.

d.   Click **OK**.

Repeat for additional groups.

Step 6   After you add the groups you want to monitor, click **Apply**.

Configuring the Cloud Web Security Policy

After you configure the ASA service policy rules, launch the ScanCenter Portal to configure Web content scanning, filtering, malware protection services, and reports.

Detailed Steps

Go to: https://scancenter.scansafe.com/portal/admin/login.jsp.

For more information, see the Cisco ScanSafe Cloud Web Security Configuration Guides:

http://www.cisco.com/en/US/products/ps11720/products_installation_and_configuration_guides_list.html

Monitoring Cloud Web Security

Command	Purpose
Monitoring > Properties > Cloud Web Security	Shows the status of the server, whether it is the current active server, the backup server, or unreachable.
	Shows total and current HTTP(S) connections. In multiple context mode, statistics are only shown within a context.
See the following URL: http://Whoami.scansafe.net	From a client, access this web site to determine if your traffic is going to the Cloud Web Security server.

Related Documents

Related Documents	URL
Cisco ScanSafe Cloud Web Security Configuration Guides	http://www.cisco.com/en/US/products/ps11720/products_installation_and_configuration_guides_list.html

Feature History for Cisco Cloud Web Security

Table 23-1 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 23-1   Feature History for Cloud Web Security

Feature Name	Platform Releases	Feature Information
Cloud Web Security	9.0(1)	This feature was introduced. Cisco Cloud Web Security provides content scanning and other malware protection service for web traffic. It can also redirect and report about web traffic based on user identity. We introduced or modified the following screens: Configuration > Device Management > Cloud Web Security Configuration > Firewall > Objects > Class Maps > Cloud Web Security Configuration > Firewall > Objects > Class Maps > Cloud Web Security > Add/Edit Configuration > Firewall > Objects > Inspect Maps > Cloud Web Security Configuration > Firewall > Objects > Inspect Maps > Cloud Web Security > Add/Edit Configuration > Firewall > Objects > Inspect Maps > Cloud Web Security > Add/Edit > Manage Cloud Web Security Class Maps Configuration > Firewall > Identity Options Configuration > Firewall > Service Policy Rules Monitoring > Properties > Cloud Web Security

Botnet Traffic Filter

Malware is malicious software that is installed on an unknowing host. Malware that attempts network activity such as sending private data (passwords, credit card numbers, key strokes, or proprietary data) can be detected by the Botnet Traffic Filter when the malware starts a connection to a known bad IP address. The Botnet Traffic Filter checks incoming and outgoing connections against a dynamic database of known bad domain names and IP addresses (the *blacklist*), and then logs or blocks any suspicious activity.

You can also supplement the Cisco dynamic database with blacklisted addresses of your choosing by adding them to a static blacklist; if the dynamic database includes blacklisted addresses that you think should not be blacklisted, you can manually enter them into a static *whitelist*. Whitelisted addresses still generate syslog messages, but because you are only targeting blacklist syslog messages, they are informational.

Note If you do not want to use the Cisco dynamic database at all, because of internal requirements, you can use the static blacklist alone if you can identify all the malware sites that you want to target.

This chapter describes how to configure the Botnet Traffic Filter and includes the following sections:

Information About the Botnet Traffic Filter

This section includes information about the Botnet Traffic Filter and includes the following topics:

Botnet Traffic Filter Address Types

Addresses monitored by the Botnet Traffic Filter include:

• Known malware addresses—These addresses are on the blacklist identified by the dynamic database and the static blacklist.

• Known allowed addresses—These addresses are on the whitelist. The whitelist is useful when an address is blacklisted by the dynamic database and also identified by the static whitelist.

• Ambiguous addresses—These addresses are associated with multiple domain names, but not all of these domain names are on the blacklist. These addresses are on the *greylist*.

• Unlisted addresses—These addresses are unknown, and not included on any list.

Botnet Traffic Filter Actions for Known Addresses

You can configure the Botnet Traffic Filter to log suspicious activity, and you can optionally configure it to block suspicious traffic automatically.

Unlisted addresses do not generate any syslog messages, but addresses on the blacklist, whitelist, and greylist generate syslog messages differentiated by type. See Botnet Traffic Filter Syslog Messaging, page 24-14 for more information.

Botnet Traffic Filter Databases

The Botnet Traffic Filter uses two databases for known addresses. You can use both databases together, or you can disable use of the dynamic database and use the static database alone. This section includes the following topics:

Information About the Dynamic Database

The Botnet Traffic Filter can receive periodic updates for the dynamic database from the Cisco update server. This database lists thousands of known bad domain names and IP addresses.

How the ASA Uses the Dynamic Database

The ASA uses the dynamic database as follows:

1. When the domain name in a DNS reply matches a name in the dynamic database, the Botnet Traffic Filter adds the name and IP address to the *DNS reverse lookup cache*.

2. When the infected host starts a connection to the IP address of the malware site, then the ASA sends a syslog message informing you of the suspicious activity and optionally drops the traffic if you configured the ASA to do so.

3. In some cases, the IP address itself is supplied in the dynamic database, and the Botnet Traffic Filter logs or drops any traffic to that IP address without having to inspect DNS requests.

Database Files

The database files are downloaded from the Cisco update server, and then stored in running memory; they are not stored in flash memory. Be sure to identify a DNS server for the ASA so that it can access the Cisco update server URL. In multiple context mode, the system downloads the database for all contexts using the admin context interface; be sure to identify a DNS server in the admin context.

If you need to delete the database, use theConfiguration > Firewall > Botnet Traffic Filter > Botnet Database pane Purge Botnet Database button instead. Be sure to first disable use of the database by unchecking the **Use Botnet data dynamically downloaded from updater server** check box in the Configuration > Firewall > Botnet Traffic Filter > Botnet Database > Dynamic Database Configuration area.

Note To filter on the domain names in the dynamic database, you need to enable DNS packet inspection with Botnet Traffic Filter snooping; the ASA looks inside the DNS packets for the domain name and associated IP address.

Database Traffic Types

The dynamic database includes the following types of addresses:

- Ads—These are advertising networks that deliver banner ads, interstitials, rich media ads, pop-ups, and pop-unders for websites, spyware and adware. Some of these networks send ad-oriented HTML emails and email verification services.

- Data Tracking—These are sources associated with companies and websites that offer data tracking and metrics services to websites and other online entities. Some of these also run small advertising networks.

- Spyware—These are sources that distribute spyware, adware, greyware, and other potentially unwanted advertising software. Some of these also run exploits to install such software.

- Malware—These are sources that use various exploits to deliver adware, spyware and other malware to victim computers. Some of these are associated with rogue online vendors and distributors of dialers which deceptively call premium-rate phone numbers.

- Adult—These are sources associated with adult networks/services offering web hosting for adult content, advertising, content aggregation, registration & billing, and age verification. These may be tied to distribution of adware, spyware, and dialers.

- Bot and Threat Networks—These are rogue systems that control infected computers. They are either systems hosted on threat networks or systems that are part of the botnet itself.

Information About the Static Database

You can manually enter domain names or IP addresses (host or subnet) that you want to tag as bad names in a blacklist. Static blacklist entries are always designated with a Very High threat level. You can also enter names or IP addresses in a whitelist, so that names or addresses that appear on both the *dynamic* blacklist and the whitelist are identified only as whitelist addresses in syslog messages and reports. Note that you see syslog messages for whitelisted addresses even if the address is not also in the dynamic blacklist.

When you add a domain name to the static database, the ASA waits 1 minute, and then sends a DNS request for that domain name and adds the domain name/IP address pairing to the *DNS host cache*. (This action is a background process, and does not affect your ability to continue configuring the ASA). We recommend also enabling DNS packet inspection with Botnet Traffic Filter snooping. The ASA uses Botnet Traffic Filter snooping instead of the regular DNS lookup to resolve static blacklist domain names in the following circumstances:

- The ASA DNS server is unavailable.

- A connection is initiated during the 1 minute waiting period before the ASA sends the regular DNS request.

If DNS snooping is used, when an infected host sends a DNS request for a name on the static database, the ASA looks inside the DNS packets for the domain name and associated IP address and adds the name and IP address to the DNS reverse lookup cache.

If you do not enable Botnet Traffic Filter snooping, and one of the above circumstances occurs, then that traffic will not be monitored by the Botnet Traffic Filter.

Information About the DNS Reverse Lookup Cache and DNS Host Cache

When you use the dynamic database with DNS snooping, entries are added to the DNS reverse lookup cache. If you use the static database, entries are added to the DNS host cache (see Information About the Static Database, page 24-3 about using the static database with DNS snooping and the DNS reverse lookup cache).

Entries in the DNS reverse lookup cache and the DNS host cache have a time to live (TTL) value provided by the DNS server. The largest TTL value allowed is 1 day (24 hours); if the DNS server provides a larger TTL, it is truncated to 1 day maximum.

For the DNS reverse lookup cache, after an entry times out, the ASA renews the entry when an infected host initiates a connection to a known address, and DNS snooping occurs.

For the DNS host cache, after an entry times out, the ASA periodically requests a refresh for the entry.

For the DNS host cache, the maximum number of blacklist entries and whitelist entries is 1000 each. The number of entries in the DNS reverse lookup cache varies per model.

How the Botnet Traffic Filter Works

Figure 24-1 shows how the Botnet Traffic Filter works with the dynamic database plus DNS inspection with Botnet Traffic Filter snooping.

Figure 24-1   How the Botnet Traffic Filter Works with the Dynamic Database

Figure 24-2 shows how the Botnet Traffic Filter works with the static database.

Figure 24-2   How the Botnet Traffic Filter Works with the Static Database

Licensing Requirements for the Botnet Traffic Filter

The following table shows the licensing requirements for this feature:

Model	License Requirement
ASAv	Standard or Premium License.
All other models	You need the following licenses: • Botnet Traffic Filter License. • Strong Encryption (3DES/AES) License to download the dynamic database.

Prerequisites for the Botnet Traffic Filter

To use the dynamic database, identify a DNS server for the ASA so that it can access the Cisco update server URL. In multiple context mode, the system downloads the database for all contexts using the admin context interface; be sure to identify a DNS server in the admin context.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Failover Guidelines

Does not support replication of the DNS reverse lookup cache, DNS host cache, or the dynamic database in Stateful Failover.

IPv6 Guidelines

Does not support IPv6.

Additional Guidelines and Limitations

- TCP DNS traffic is not supported.
- You can add up to 1000 blacklist entries and 1000 whitelist entries in the static database.
- The packet tracer is not supported.

Default Settings

By default, the Botnet Traffic Filter is disabled, as is use of the dynamic database.

For DNS inspection, which is enabled by default, Botnet Traffic Filter snooping is disabled by default.

Configuring the Botnet Traffic Filter

This section includes the following topics:

Task Flow for Configuring the Botnet Traffic Filter

To configure the Botnet Traffic Filter, perform the following steps:

Step 1   Enable use of the dynamic database. See Configuring the Dynamic Database, page 24-8.

This procedure enables database updates from the Cisco update server, and also enables use of the downloaded dynamic database by the ASA. Disallowing use of the downloaded database is useful in multiple context mode so you can configure use of the database on a per-context basis.

Step 2   (Optional) Add static entries to the database. See Adding Entries to the Static Database, page 24-9.

This procedure lets you augment the dynamic database with domain names or IP addresses that you want to blacklist or whitelist. You might want to use the static database instead of the dynamic database if you do not want to download the dynamic database over the Internet.

Step 3   Enable DNS snooping. See Enabling DNS Snooping, page 24-9.

This procedure enables inspection of DNS packets, compares the domain name with those in the dynamic database or the static database (when a DNS server for the ASA is unavailable), and adds the name and IP address to the DNS reverse lookup cache. This cache is then used by the Botnet Traffic Filter when connections are made to the suspicious address.

Step 4   Enable traffic classification and actions for the Botnet Traffic Filter. See Enabling Traffic Classification and Actions for the Botnet Traffic Filter, page 24-10.

This procedure enables the Botnet Traffic Filter, which compares the source and destination IP address in each initial connection packet to the IP addresses in the dynamic database, static database, DNS reverse lookup cache, and DNS host cache, and sends a syslog message or drops any matching traffic.

Step 5   (Optional) Block traffic manually based on syslog message information. See Blocking Botnet Traffic Manually, page 24-12.

If you choose not to block malware traffic automatically, you can block traffic manually by configuring an access rule to deny traffic, or by using the **shun** command in the Command Line Interface tool to block all traffic to and from a host.

Configuring the Dynamic Database

This procedure enables database updates, and also enables use of the downloaded dynamic database by the ASA. In multiple context mode, the system downloads the database for all contexts using the admin context interface. You can configure *use* of the database on a per-context basis.

By default, downloading and using the dynamic database is disabled.

Prerequisites

Enable ASA use of a DNS server in the Device Management > DNS > DNS Client > DNS Lookup area. In multiple context mode, the system downloads the database for all contexts using the admin context interface; be sure to identify a DNS server in the admin context.

Detailed Steps

Step 1 Enable downloading of the dynamic database.

- In Single mode, choose the **Configuration > Firewall > Botnet Traffic Filter > Botnet Database** pane, then check the **Enable Botnet Updater Client** check box.

- In multiple context mode in the System execution space, choose the **Configuration > Device Management > Botnet Database** pane, then check the **Enable Botnet Updater Client** check box.

This setting enables downloading of the dynamic database from the Cisco update server. In multiple context mode, enter this command in the system execution space. If you do not have a database already installed on the ASA, it downloads the database after approximately 2 minutes. The update server determines how often the ASA polls the server for future updates, typically every hour.

Step 2 (Multiple context mode only) In multiple context mode, click **Apply**. Then change to the context where you want to configure the Botnet Traffic Filter by double-clicking the context name in the Device List.

Step 3 In the Configuration > Firewall > Botnet Traffic Filter > Botnet Database > Dynamic Database Configuration area, check the **Use Botnet data dynamically downloaded from updater server** check box.

Step 4 Click **Apply**.

Step 5 (Optional) If you want to later remove the database from running memory, perform the following steps:

   a. Disable use of the database by unchecking the **Use Botnet data dynamically downloaded from updater server** check box.

   b. Click **Apply**.

   c. Click **Purge Botnet Database**.

   d. To redownload the database, re-check the **Use Botnet data dynamically downloaded from updater server** check box.

   e. Click **Apply**.

Note The Fetch Botnet Database button is for testing purposes only; it downloads and verifies the dynamic database, but does not store it in running memory.

For information about the Search Dynamic Database area, see Searching the Dynamic Database, page 24-13.

What to Do Next

See Adding Entries to the Static Database, page 24-9.

Adding Entries to the Static Database

The static database lets you augment the dynamic database with domain names or IP addresses that you want to blacklist or whitelist. Static blacklist entries are always designated with a Very High threat level. See Information About the Static Database, page 24-3 for more information.

Prerequisites

- In multiple context mode, perform this procedure in the context execution space.
- Enable ASA use of a DNS server in the Device Management > DNS > DNS Client > DNS Lookup area. In multiple context mode, enable DNS per context.

Detailed Steps

Step 1 Choose the **Configuration > Firewall > Botnet Traffic Filter > Black** or White List pane, click **Add** for the Whitelist or Blacklist.

The Enter hostname or IP Address dialog box appears.

Step 2 In the Addresses field, enter one or more domain names, IP addresses, and IP address/netmasks.

Enter multiple entries separated by commas, spaces, lines, or semi-colons. You can enter up to 1000 entries for each type.

Step 3 Click **OK**.

Step 4 Click **Apply**.

What to Do Next

See Enabling DNS Snooping, page 24-9.

Enabling DNS Snooping

This procedure enables inspection of DNS packets and enables Botnet Traffic Filter snooping, which compares the domain name with those on the dynamic database or static database, and adds the name and IP address to the Botnet Traffic Filter DNS reverse lookup cache. This cache is then used by the Botnet Traffic Filter when connections are made to the suspicious address.

Prerequisites

- In multiple context mode, perform this procedure in the context execution space.

- You must first configure DNS inspection for traffic that you want to snoop using the Botnet Traffic Filter. See DNS Inspection, page 9-1 and Chapter 1, "Service Policy," for detailed information about configuring advanced DNS inspection options using the Modular Policy Framework.

> ✎
>
> **Note**   You can also configure DNS snooping directly in the Configuration > Firewall > Service Policy Rules > Rule Actions > Protocol Inspection > Select DNS Inspect Map dialog box by checking the **Enable Botnet traffic filter DNS snooping** check box.

Restrictions

TCP DNS traffic is not supported.

Default DNS Inspection Configuration and Recommended Configuration

The default configuration for DNS inspection inspects all UDP DNS traffic on all interfaces, and does not have DNS snooping enabled.

We suggest that you enable DNS snooping only on interfaces where external DNS requests are going. Enabling DNS snooping on all UDP DNS traffic, including that going to an internal DNS server, creates unnecessary load on the ASA.

For example, if the DNS server is on the outside interface, you should enable DNS inspection with snooping for all UDP DNS traffic on the outside interface.

Detailed Steps

Step 1   Choose the **Configuration > Firewall > Botnet Traffic Filter > DNS Snooping** pane.

All existing service rules that include DNS inspection are listed in the table.

Step 2   For each rule for which you want to enable DNS snooping, in the DNS Snooping Enabled column, check the check box.

Step 3   Click **Apply**.

What to Do Next

See Enabling Traffic Classification and Actions for the Botnet Traffic Filter, page 24-10.

Enabling Traffic Classification and Actions for the Botnet Traffic Filter

This procedure enables the Botnet Traffic Filter. The Botnet Traffic Filter compares the source and destination IP address in each initial connection packet to the following:

- Dynamic database IP addresses
- Static database IP addresses
- DNS reverse lookup cache (for dynamic database domain names)

- DNS host cache (for static database domain names)

When an address matches, the ASA sends a syslog message. The only additional action currently available is to drop the connection.

Prerequisites

In multiple context mode, perform this procedure in the context execution space.

Recommended Configuration

Although DNS snooping is not required, we recommend configuring DNS snooping for maximum use of the Botnet Traffic Filter (see Enabling DNS Snooping, page 24-9). Without DNS snooping for the dynamic database, the Botnet Traffic Filter uses only the static database entries, plus any IP addresses in the dynamic database; domain names in the dynamic database are not used.

We recommend enabling the Botnet Traffic Filter on all traffic on the Internet-facing interface, and enabling dropping of traffic with a severity of moderate and higher.

Detailed Steps

Step 1   Choose the **Configuration > Firewall > Botnet Traffic Filter > Traffic Settings** pane.

Step 2   To enable the Botnet Traffic Filter on specified traffic, perform the following steps:

a.   In the Traffic Classification area, check the **Traffic Classified** check box for each interface on which you want to enable the Botnet Traffic Filter.

You can configure a global classification that applies to all interfaces by checking the Traffic Classified check box for Global (All Interfaces). If you configure an interface-specific classification, the settings for that interface overrides the global setting.

b.   For each interface, from the **ACL Used** drop-down list choose either --ALL TRAFFIC-- (the default), or any ACL configured on the ASA.

For example, you might want to monitor all port 80 traffic on the outside interface.

To add or edit ACLs, click **Manage ACL** to bring up the ACL Manager. See the general operations configuration guide for more information.

Step 3   (Optional) To treat greylisted traffic as blacklisted traffic for action purposes, in the Ambiguous Traffic Handling area, check the **Treat ambiguous (greylisted) traffic as malicious (blacklisted) traffic** check box.

If you do not enable this option, greylisted traffic will not be dropped if you configure a rule in the Blacklisted Traffic Actions area. See Botnet Traffic Filter Address Types, page 24-2 for more information about the greylist.

Step 4   (Optional) To automatically drop malware traffic, perform the following steps.

To manually drop traffic, see Blocking Botnet Traffic Manually, page 24-12.

a.   In the Blacklisted Traffic Actions area, click **Add**.

The Add Blacklisted Traffic Action dialog box appears.

b.   From the Interface drop-down list, choose the interface on which you want to drop traffic. Only interfaces on which you enabled Botnet Traffic Filter traffic classification are available.

c. In the Threat Level area, choose one of the following options to drop traffic specific threat levels. The default level is a range between Moderate and Very High.

> ✎
>
> **Note**   We highly recommend using the default setting unless you have strong reasons for changing the setting.

- Value—Specify the threat level you want to drop:
    - **Very Low**
    - **Low**
    - **Moderate**
    - **High**
    - **Very High**

    > ✎
    >
    > **Note**   Static blacklist entries are always designated with a Very High threat level.

- Range—Specify a range of threat levels.

d. In the ACL Used area, from the **ACL Used** drop-down list choose either --ALL TRAFFIC-- (the default), or any ACL configured on the ASA.

> ✎
>
> **Note**   Be sure the ACL is a subset of the traffic you specified in the Traffic Classification area.

To add or edit ACLs, click **Manage** to bring up the ACL Manager. See general operations configuration guide for more information.

e. Click **OK**.

You return to the Traffic Settings pane.

f. If you want to apply additional rules to a given interface, repeat steps a through e.

Make sure you do not specify overlapping traffic in multiple rules for a given interface. Because you cannot control the exact order that rules are matched, overlapping traffic means you do not know which command will be matched. For example, do not specify both a rule that matches --ALL TRAFFIC-- as well as a command with and ACL for a given interface. In this case, the traffic might never match the command with the ACL. Similarly, if you specify multiple commands with ACLs, make sure each ACL is unique, and that the networks do not overlap.

Step 5   Click **Apply**.

Blocking Botnet Traffic Manually

If you choose not to block malware traffic automatically (see Enabling Traffic Classification and Actions for the Botnet Traffic Filter, page 24-10), you can block traffic manually by configuring an access rule to deny traffic, or by using the **shun** command in the Command Line Interface tool to block all traffic to and from a host. For some messages, you can automatically configure access rules in ASDM.

For example, you receive the following syslog message:

```
ASA-4-338002: Dynamic Filter permitted black listed TCP traffic from inside:10.1.1.45/6798
(209.165.201.1/7890) to outside:209.165.202.129/80 (209.165.202.129/80), destination
209.165.202.129 resolved from dynamic list: bad.example.com
```

You can then perform one of the following actions:

- Create an access rule to deny traffic.

  For example, using the syslog message above, you might want to deny traffic from the infected host at 10.1.1.45 to the malware site at 209.165.202.129. Or, if there are many connections to different blacklisted addresses, you can create an ACL to deny all traffic from 10.1.1.45 until you resolve the infection on the host computer.

  For the following syslog messages, a reverse access rule can be automatically created from the Real Time Log Viewer:

    – 338001, 338002, 338003, 338004 (blacklist)

    – 338201, 338202 (greylist)

  See the general operations configuration guide and Chapter 3, "Access Rules," for more information about creating an access rule.

  > **Note**   If you create a reverse access rule form a Botnet Traffic Filter syslog message, and you do not have any other access rules applied to the interface, then you might inadvertently block all traffic. Normally, without an access rule, all traffic from a high security to a low security interface is allowed. But when you apply an access rule, all traffic is denied except traffic that you explicitly permit. Because the reverse access rule is a deny rule, be sure to edit the resulting access policy for the interface to permit other traffic.
  >
  > ACLs block all future connections. To block the current connection, if it is still active, enter the **clear conn** command. For example, to clear only the connection listed in the syslog message, enter the **clear conn address 10.1.1.45 address 209.165.202.129** command. See the command reference for more information.

- Shun the infected host.

  Shunning blocks all connections from the host, so you should use an ACL if you want to block connections to certain destination addresses and ports. To shun a host, enter the following command in Tools > Command Line Interface. To drop the current connection as well as blocking all future connections, enter the destination address, source port, destination port, and optional protocol.

  shun *src_ip* [*dst_ip src_port dest_port* [*protocol*]]

  For example, to block future connections from 10.1.1.45, and also drop the current connection to the malware site in the syslog message, enter:

  shun 10.1.1.45 209.165.202.129 6798 80

After you resolve the infection, be sure to remove the ACL or the shun. To remove the shun, enter **no shun** *src_ip*.

Searching the Dynamic Database

If you want to check if a domain name or IP address is included in the dynamic database, you can search the database for a string.

Detailed Steps

Step 1   Go to the Search Dynamic Database area:

- In Single mode or within a context, choose the **Configuration > Firewall > Botnet Traffic Filter > Botnet Database Update** pane.

- In multiple context mode in the System execution space, choose the **Configuration > Device Management > Botnet Database Update** pane.

Step 2   In the Search string field, enter a string at least 3 characters in length, and click **Find Now**.

The first two matches are shown. To refine your search for a more specific match, enter a longer string.

Step 3   To clear the displayed matches and the search string, click **Clear**, or you can just enter a new string and click **Find Now** to get a new display.

Monitoring the Botnet Traffic Filter

Whenever a known address is classified by the Botnet Traffic Filter, then a syslog message is generated. You can also monitor Botnet Traffic Filter statistics and other parameters by entering commands on the ASA. This section includes the following topics:

Botnet Traffic Filter Syslog Messaging

The Botnet Traffic Filter generates detailed syslog messages numbered 338*nnn*. Messages differentiate between incoming and outgoing connections, blacklist, whitelist, or greylist addresses, and many other variables. (The greylist includes addresses that are associated with multiple domain names, but not all of these domain names are on the blacklist.)

See the syslog messages guide for detailed information about syslog messages.

For the following syslog messages, a reverse access rule can be automatically created from the Real Time Log Viewer:

- 338001, 338002, 338003, 338004 (blacklist)

- 338201, 338202 (greylist)

Botnet Traffic Filter Monitor Panes

To monitor the Botnet Traffic Filter, see the following panes:

Command	Purpose
Home > Firewall Dashboard	Shows the Top Botnet Traffic Filter Hits, which shows reports of the top 10 malware sites, ports, and infected hosts. This report is a snapshot of the data, and may not match the top 10 items since the statistics started to be collected. If you right-click an IP address, you can invoke the whois tool to learn more about the botnet site. • Top Malware Sites—Shows top malware sites. • Top Malware Ports—Shows top malware ports. • Top Infected Hosts—Shows the top infected hosts.
Monitoring > Botnet Traffic Filter > Statistics	Shows how many connections were classified as whitelist, blacklist, and greylist connections, and how many connections were dropped. (The greylist includes addresses that are associated with multiple domain names, but not all of these domain names are on the blacklist.) The Details button shows how many packets at each threat level were classified or dropped.
Monitoring > Botnet Traffic Filter > Real-time Reports	Generates reports of the top 10 malware sites, ports, and infected hosts monitored. The top 10 malware-sites report includes the number of connections dropped, and the threat level and category of each site. This report is a snapshot of the data, and may not match the top 10 items since the statistics started to be collected. If you right-click a site IP address, you can invoke the whois tool to learn more about the malware site. Reports can be saved as a PDF file.
Monitoring > Botnet Traffic Filter > Infected Hosts	Generates reports about infected hosts. These reports contain detailed history about infected hosts, showing the correlation between infected hosts, visited malware sites, and malware ports. The Maximum Connections option shows the 20 infected hosts with the most number of connections. The Latest Activity option shows the 20 hosts with the most recent activity. The Highest Threat Level option shows the 20 hosts that connected to the malware sites with the highest threat level. The Subnet option shows up to 20 hosts within the specified subnet. Reports can be saved as a PDF file, as either the Current View or the Whole Buffer. The Whole Buffer option shows all buffered infected-hosts information.
Monitoring > Botnet Traffic Filter > Updater Client	Shows information about the updater server, including the server IP address, the next time the ASA will connect with the server, and the database version last installed.
Monitoring > Botnet Traffic Filter > DNS Snooping	Shows the Botnet Traffic Filter DNS snooping actual IP addresses and names. All inspected DNS data is included in this output, and not just matching names in the blacklist. DNS data from static entries are not included.

Command	Purpose
Monitoring > Botnet Traffic Filter > Dynamic Database	Shows information about the dynamic database, including when the dynamic database was last downloaded, the version of the database, how many entries the database contains, and 10 sample entries.
Monitoring > Botnet Traffic Filter > ASP Table Hits	Shows the Botnet Traffic Filter rules that are installed in the accelerated security path.

Where to Go Next

- To configure the syslog server, see the general operations configuration guide.
- To block connections with an access rule, see Chapter 3, "Access Rules."

Feature History for the Botnet Traffic Filter

Table 24-1 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 24-1   Feature History for the Botnet Traffic Filter

Feature Name	Platform Releases	Feature Information
Botnet Traffic Filter	8.2(1)	This feature was introduced.
Automatic blocking, and blacklist category and threat level reporting.	8.2(2)	The Botnet Traffic Filter now supports automatic blocking of blacklisted traffic based on the threat level. You can also view the category and threat level of malware sites in statistics and reports. The 1 hour timeout for reports for top hosts was removed; there is now no timeout. The following screens were introduced or modified: Configuration > Firewall > Botnet Traffic Filter > Traffic Settings, and Monitoring > Botnet Traffic Filter > Infected Hosts.

Threat Detection

This chapter describes how to configure threat detection statistics and scanning threat detection and includes the following sections:

Information About Threat Detection

The threat detection feature consists of the following elements:

- Different levels of statistics gathering for various threats.

  Threat detection statistics can help you manage threats to your ASA; for example, if you enable scanning threat detection, then viewing statistics can help you analyze the threat. You can configure two types of threat detection statistics:

  - Basic threat detection statistics—Includes information about attack activity for the system as a whole. Basic threat detection statistics are enabled by default and have no performance impact.

  - Advanced threat detection statistics—Tracks activity at an object level, so the ASA can report activity for individual hosts, ports, protocols, or ACLs. Advanced threat detection statistics can have a major performance impact, depending on the statistics gathered, so only the ACL statistics are enabled by default.

- Scanning threat detection, which determines when a host is performing a scan.

  You can optionally shun any hosts determined to be a scanning threat.

Licensing Requirements for Threat Detection

The following table shows the licensing requirements for this feature:

Model	License Requirement
ASAv	Standard or Premium License.
All other models	Base License.

Configuring Basic Threat Detection Statistics

Basic threat detection statistics include activity that might be related to an attack, such as a DoS attack.

This section includes the following topics:

- Information About Basic Threat Detection Statistics, page 25-2
- Guidelines and Limitations, page 25-3
- Default Settings, page 25-3
- Configuring Basic Threat Detection Statistics, page 25-4
- Monitoring Basic Threat Detection Statistics, page 25-4
- Feature History for Basic Threat Detection Statistics, page 25-5

Information About Basic Threat Detection Statistics

Using basic threat detection statistics, the ASA monitors the rate of dropped packets and security events due to the following reasons:

- Denial by ACLs
- Bad packet format (such as invalid-ip-header or invalid-tcp-hdr-length)
- Connection limits exceeded (both system-wide resource limits, and limits set in the configuration)
- DoS attack detected (such as an invalid SPI, Stateful Firewall check failure)
- Basic firewall checks failed (This option is a combined rate that includes all firewall-related packet drops in this bulleted list. It does not include non-firewall-related drops such as interface overload, packets failed at application inspection, and scanning attack detected.)
- Suspicious ICMP packets detected
- Packets failed application inspection
- Interface overload
- Scanning attack detected (This option monitors scanning attacks; for example, the first TCP packet is not a SYN packet, or the TCP connection failed the 3-way handshake. Full scanning threat detection (see Configuring Scanning Threat Detection, page 25-8) takes this scanning attack rate information and acts on it by classifying hosts as attackers and automatically shunning them, for example.)
- Incomplete session detection such as TCP SYN attack detected or no data UDP session attack detected

When the ASA detects a threat, it immediately sends a system log message (733100). The ASA tracks two types of rates: the average event rate over an interval, and the burst event rate over a shorter burst interval. The burst rate interval is 1/30th of the average rate interval or 10 seconds, whichever is higher.

For each received event, the ASA checks the average and burst rate limits; if both rates are exceeded, then the ASA sends two separate system messages, with a maximum of one message for each rate type per burst period.

Basic threat detection affects performance only when there are drops or potential threats; even in this scenario, the performance impact is insignificant.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature:

Security Context Guidelines
Supported in single mode only. Multiple mode is not supported.

Firewall Mode Guidelines
Supported in routed and transparent firewall mode.

Types of Traffic Monitored
Only through-the-box traffic is monitored; to-the-box traffic is not included in threat detection.

Default Settings

Basic threat detection statistics are enabled by default.

Table 25-1 lists the default settings. You can view all these default settings using the **show running-config all threat-detection** command in Tools > Command Line Interface.

Table 25-1   Basic Threat Detection Default Settings

Packet Drop Reason	Trigger Settings	
	Average Rate	Burst Rate
• DoS attack detected • Bad packet format • Connection limits exceeded • Suspicious ICMP packets detected	100 drops/sec over the last 600 seconds.	400 drops/sec over the last 20 second period.
	80 drops/sec over the last 3600 seconds.	320 drops/sec over the last 120 second period.
Scanning attack detected	5 drops/sec over the last 600 seconds.	10 drops/sec over the last 20 second period.
	4 drops/sec over the last 3600 seconds.	8 drops/sec over the last 120 second period.
Incomplete session detected such as TCP SYN attack detected or no data UDP session attack detected (combined)	100 drops/sec over the last 600 seconds.	200 drops/sec over the last 20 second period.
	80 drops/sec over the last 3600 seconds.	160 drops/sec over the last 120 second period.

Table 25-1   Basic Threat Detection Default Settings (continued)

Packet Drop Reason	Trigger Settings	
	Average Rate	Burst Rate
Denial by ACLs	400 drops/sec over the last 600 seconds.	800 drops/sec over the last 20 second period.
	320 drops/sec over the last 3600 seconds.	640 drops/sec over the last 120 second period.
• Basic firewall checks failed • Packets failed application inspection	400 drops/sec over the last 600 seconds.	1600 drops/sec over the last 20 second period.
	320 drops/sec over the last 3600 seconds.	1280 drops/sec over the last 120 second period.
Interface overload	2000 drops/sec over the last 600 seconds.	8000 drops/sec over the last 20 second period.
	1600 drops/sec over the last 3600 seconds.	6400 drops/sec over the last 120 second period.

Configuring Basic Threat Detection Statistics

This section describes how to configure basic threat detection statistics, including enabling or disabling it and changing the default limits.

Detailed Steps

Step 1   To enable or disable basic threat detection, choose the **Configuration > Firewall > Threat Detection** pane, and check the **Enable Basic Threat Detection** check box.

Step 2   Click **Apply**.

Monitoring Basic Threat Detection Statistics

To monitor basic threat detection statistics, perform the following task:

Path	Purpose
Home > Firewall Dashboard > Traffic Overview	Displays basic threat detection statistics. For a description of each event type, see Information About Basic Threat Detection Statistics, page 25-2.

Feature History for Basic Threat Detection Statistics

Table 25-2 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 25-2   Feature History for Basic Threat Detection Statistics

Feature Name	Platform Releases	Feature Information
Basic threat detection statistics	8.0(2)	Basic threat detection statistics was introduced. The following screen was introduced: Configuration > Firewall > Threat Detection, Home > Firewall Dashboard > Traffic Overview.
Burst rate interval changed to 1/30th of the average rate.	8.2(1)	In earlier releases, the burst rate interval was 1/60th of the average rate. To maximize memory usage, the sampling interval was reduced to 30 times during the average rate.
Improved memory usage	8.3(1)	The memory usage for threat detection was improved.

Configuring Advanced Threat Detection Statistics

You can configure the ASA to collect extensive statistics. This section includes the following topics:

- Information About Advanced Threat Detection Statistics, page 25-5
- Guidelines and Limitations, page 25-5
- Default Settings, page 25-6
- Configuring Advanced Threat Detection Statistics, page 25-6
- Monitoring Advanced Threat Detection Statistics, page 25-7
- Feature History for Advanced Threat Detection Statistics, page 25-8

Information About Advanced Threat Detection Statistics

Advanced threat detection statistics show both allowed and dropped traffic rates for individual objects such as hosts, ports, protocols, or ACLs.

Caution   Enabling advanced statistics can affect the ASA performance, depending on the type of statistics enabled. Enabling host statistics affects performance in a significant way; if you have a high traffic load, you might consider enabling this type of statistics temporarily. Port statistics, however, has modest impact.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature:

Security Context Guidelines

Only TCP Intercept statistics are available in multiple mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Types of Traffic Monitored

Only through-the-box traffic is monitored; to-the-box traffic is not included in threat detection.

Default Settings

By default, statistics for ACLs are enabled.

Configuring Advanced Threat Detection Statistics

By default, statistics for ACLs are enabled. To enable other statistics, perform the following steps.

Detailed Steps

Step 1   Choose the **Configuration > Firewall > Threat Detection** pane.

Step 2   In the Scanning Threat Statistics area, choose one of the following options:

- Enable *all* statistics—Click the **Enable All Statistics** radio button.
- Disable *all* statistics—Click the **Disable All Statistics** radio button.
- Enable only certain statistics—Click the **Enable Only Following Statistics** radio button.

Step 3   If you chose to **Enable Only Following Statistics**, then check one or more of the following check boxes:

- **Hosts**—Enables host statistics. The host statistics accumulate for as long as the host is active and in the scanning threat host database. The host is deleted from the database (and the statistics cleared) after 10 minutes of inactivity.
- **Access Rules** (enabled by default)—Enables statistics for access rules.
- **Port**—Enables statistics for TCP and UDP ports.
- **Protocol**—Enables statistics for non-TCP/UDP IP protocols.
- **TCP-Intercept**—Enables statistics for attacks intercepted by TCP Intercept (see Configuring Connection Settings, page 20-8 to enable TCP Intercept).

Step 4   For host, port, and protocol statistics, you can change the number of rate intervals collected. In the Rate Intervals area, choose **1 hour**, **1 and 8 hours**, or **1, 8 and 24 hours** for each statistics type. The default interval is **1 hour**, which keeps the memory usage low.

Step 5   For TCP Intercept statistics, you can set the following options in the TCP Intercept Threat Detection area:

- **Monitoring Window Size**—Sets the size of the history monitoring window, between 1 and 1440 minutes. The default is 30 minutes. The ASA samples the number of attacks 30 times during the rate interval, so for the default 30 minute period, statistics are collected every 60 seconds.

- **Burst Threshold Rate**—Sets the threshold for syslog message generation, between 25 and 2147483647. The default is 400 per second. When the burst rate is exceeded, syslog message 733104 is generated.

- **Average Threshold Rate**—Sets the average rate threshold for syslog message generation, between 25 and 2147483647. The default is 200 per second. When the average rate is exceeded, syslog message 733105 is generated.

Click **Set Default** to restore the default values.

Step 6   Click **Apply**.

Monitoring Advanced Threat Detection Statistics

To monitor advanced threat detection statistics, perform one of the following tasks:

Path	Purpose
Home > Firewall Dashboard > Top 10 Access Rules Home > Firewall Dashboard > Top Usage Statistics	Displays the top 10 statistics.
	For the Top 10 Access Rules, permitted and denied traffic are not differentiated in this display. In the Traffic Overview > Dropped Packets Rate graph, you can track ACL denies.
	The Top 10 Sources and Top 10 Destinations tabs show statistics for hosts. **Note**: Due to the threat detction algorithm, an interface used as a combination failover and state link could appear in the top 10 hosts; this is expected behavior, and you can ignore this IP address in the display.
	The Top 10 Services tab shows statistics for both ports and protocols (both must be enabled for the display), and shows the combined statistics of TCP/UDP port and IP protocol types. TCP (protocol 6) and UDP (protocol 17) are not included in the display for IP protocols; TCP and UDP ports are, however, included in the display for ports. If you only enable statistics for one of these types, port or protocol, then you will only view the enabled statistics.
	The Top Ten Protected Servers under SYN Attack area shows the TCP Intercept statistics. The display includes the top 10 protected servers under attack. The **detail** button shows history sampling data. The ASA samples the number of attacks 30 times during the rate interval, so for the default 30 minute period, statistics are collected every 60 seconds.
	From the Interval drop-down list, choose **Last 1 hour**, **Last 8 hour**, or **Last 24 hour**.

Feature History for Advanced Threat Detection Statistics

Table 25-3 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 25-3   Feature History for Advanced Threat Detection Statistics

Feature Name	Platform Releases	Feature Information
Advanced threat detection statistics	8.0(2)	Advanced threat detection statistics was introduced. The following screens were introduced: Configuration > Firewall > Threat Detection, Home > Firewall Dashboard > Top 10 Access Rules, Home > Firewall Dashboard > Top Usage Status, Home > Firewall Dashboard > Top 10 Protected Servers Under SYN Attack.
TCP Intercept statistics	8.0(4)/8.1(2)	TCP Intercept statistics were introduced. The following screens were introduced or modified: Configuration > Firewall > Threat Detection, Home > Firewall Dashboard > Top 10 Protected Servers Under SYN Attack.
Customize host statistics rate intervals	8.1(2)	You can now customize the number of rate intervals for which statistics are collected. The default number of rates was changed from 3 to 1. The following screen was modified: Configuration > Firewall > Threat Detection.
Burst rate interval changed to 1/30th of the average rate.	8.2(1)	In earlier releases, the burst rate interval was 1/60th of the average rate. To maximize memory usage, the sampling interval was reduced to 30 times during the average rate.
Customize port and protocol statistics rate intervals	8.3(1)	You can now customize the number of rate intervals for which statistics are collected. The default number of rates was changed from 3 to 1. The following screen was modified: Configuration > Firewall > Threat Detection.
Improved memory usage	8.3(1)	The memory usage for threat detection was improved.

Configuring Scanning Threat Detection

This section includes the following topics:

Information About Scanning Threat Detection

A typical scanning attack consists of a host that tests the accessibility of every IP address in a subnet (by scanning through many hosts in the subnet or sweeping through many ports in a host or subnet). The scanning threat detection feature determines when a host is performing a scan. Unlike IPS scan detection that is based on traffic signatures, the ASA scanning threat detection feature maintains an extensive database that contains host statistics that can be analyzed for scanning activity.

The host database tracks suspicious activity such as connections with no return activity, access of closed service ports, vulnerable TCP behaviors such as non-random IPID, and many more behaviors.

If the scanning threat rate is exceeded, then the ASA sends a syslog message (733101), and optionally shuns the attacker. The ASA tracks two types of rates: the average event rate over an interval, and the burst event rate over a shorter burst interval. The burst event rate is 1/30th of the average rate interval or 10 seconds, whichever is higher. For each event detected that is considered to be part of a scanning attack, the ASA checks the average and burst rate limits. If either rate is exceeded for traffic sent from a host, then that host is considered to be an attacker. If either rate is exceeded for traffic received by a host, then that host is considered to be a target.

⚠
Caution   The scanning threat detection feature can affect the ASA performance and memory significantly while it creates and gathers host- and subnet-based data structure and information.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature:

Security Context Guidelines

Supported in single mode only. Multiple mode is not supported.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Types of Traffic Monitored

• Only through-the-box traffic is monitored; to-the-box traffic is not included in threat detection.

• Traffic that is denied by an ACL does not trigger scanning threat detection; only traffic that is allowed through the ASA and that creates a flow is affected by scanning threat detection.

Default Settings

Table 25-4 lists the default rate limits for scanning threat detection.

Table 25-4   Default Rate Limits for Scanning Threat Detection

Average Rate	Burst Rate
5 drops/sec over the last 600 seconds.	10 drops/sec over the last 20 second period.
5 drops/sec over the last 3600 seconds.	10 drops/sec over the last 120 second period.

The burst rate is calculated as the average rate every *N* seconds, where *N* is the burst rate interval. The burst rate interval is 1/30th of the rate interval or 10 seconds, whichever is larger.

Configuring Scanning Threat Detection

Detailed Steps

Step 1   Choose the **Configuration > Firewall > Threat Detection** pane, and check the **Enable Scanning Threat Detection** check box.

Step 2   (Optional) To automatically terminate a host connection when the ASA identifies the host as an attacker, check the **Shun Hosts detected by scanning threat** check box.

Step 3   (Optional) To except host IP addresses from being shunned, enter an address in the Networks excluded from shun field.

You can enter multiple addresses or subnets separated by commas. To choose a network from the list of IP address objects, click the **...** button.

Step 4   (Optional) To set the duration of a shun for an attacking host, check the **Set Shun Duration** check box and enter a value between 10 and 2592000 seconds. The default length is 3600 seconds (1 hour). To restore the default value, click **Set Default**.

Feature History for Scanning Threat Detection

Table 25-5 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 25-5   Feature History for Scanning Threat Detection

Feature Name	Platform Releases	Feature Information
Scanning threat detection	8.0(2)	Scanning threat detection was introduced. The following screen was introduced: Configuration > Firewall > Threat Detection.
Shun duration	8.0(4)/8.1(2)	You can now set the shun duration, The following screen was modified: Configuration > Firewall > Threat Detection.
Burst rate interval changed to 1/30th of the average rate.	8.2(1)	In earlier releases, the burst rate interval was 1/60th of the average rate. To maximize memory usage, the sampling interval was reduced to 30 times during the average rate.
Improved memory usage	8.3(1)	The memory usage for threat detection was improved.

P ART 7

ASA Modules

CHAPTER **26**

ASA FirePOWER (SFR) Module

This chapter describes how to configure the ASA FirePOWER module that runs on the ASA.

The ASA FirePOWER Module

The ASA FirePOWER module supplies next-generation firewall services, including Next-Generation IPS (NGIPS), Application Visibility and Control (AVC), URL filtering, and Advanced Malware Protection (AMP).You can use the module in single or multiple context mode, and in routed or transparent mode.

The module is also known as ASA SFR.

Although the module has a basic command line interface (CLI) for initial configuration and troubleshooting, you configure the security policy on the device using a separate application, FireSIGHT Management Center, which can be hosted on a separate FireSIGHT Management Center appliance or as a virtual appliance running on a VMware server. (FireSIGHT Management Center is also known as Defense Center.)

How the ASA FirePOWER Module Works with the ASA

You can configure your ASA FirePOWER module using one of the following deployment models:

- Inline mode—In an inline deployment, the actual traffic is sent to the ASA FirePOWER module, and the module's policy affects what happens to the traffic. After dropping undesired traffic and taking any other actions applied by policy, the traffic is returned to the ASA for further processing and ultimate transmission.

- Inline tap monitor-only mode (ASA inline)—In an inline tap monitor-only deployment, a copy of the traffic is sent to the ASA FirePOWER module, but it is not returned to the ASA. Inline tap mode lets you see what the ASA FirePOWER module would have done to traffic, and lets you evaluate the content of the traffic, without impacting the network. However, in this mode, the ASA does apply its policies to the traffic, so traffic can be dropped due to access rules, TCP normalization, and so forth.

Be sure to configure consistent policies on the ASA and the ASA FirePOWER. Both policies should reflect the inline or monitor-only mode of the traffic.

The following sections explain these modes in more detail.

ASA FirePOWER Inline Mode

In inline mode, traffic goes through the firewall checks before being forwarded to the ASA FirePOWER module. When you identify traffic for ASA FirePOWER inspection on the ASA, traffic flows through the ASA and the module as follows:

1. Traffic enters the ASA.

2. Incoming VPN traffic is decrypted.

3. Firewall policies are applied.

4. Traffic is sent to the ASA FirePOWER module.

5. The ASA FirePOWER module applies its security policy to the traffic, and takes appropriate actions.

6. Valid traffic is sent back to the ASA; the ASA FirePOWER module might block some traffic according to its security policy, and that traffic is not passed on.

7. Outgoing VPN traffic is encrypted.

8. Traffic exits the ASA.

The following figure shows the traffic flow when using the ASA FirePOWER module in inline mode. In this example, the module blocks traffic that is not allowed for a certain application. All other traffic is forwarded through the ASA.

Figure 26-1   ASA FirePOWER Module Traffic Flow in the ASA
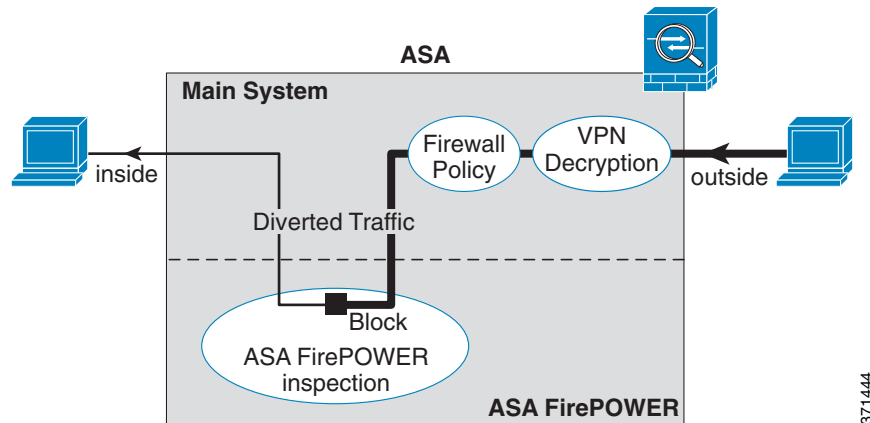
> **Note** If you have a connection between hosts on two ASA interfaces, and the ASA FirePOWER service policy is only configured for one of the interfaces, then all traffic between these hosts is sent to the ASA FirePOWER module, including traffic originating on the non-ASA FirePOWER interface (because the feature is bidirectional).

ASA FirePOWER Inline Tap Monitor-Only Mode

This mode sends a duplicate stream of traffic to the ASA FirePOWER module for monitoring purposes only. The module applies the security policy to the traffic and lets you know what it would have done if it were operating in inline mode; for example, traffic might be marked "would have dropped" in events. You can use this information for traffic analysis and to help you decide if inline mode is desirable.

> **Note** You cannot configure both inline tap monitor-only mode and normal inline mode at the same time on the ASA. Only one type of security policy is allowed. In multiple context mode, you cannot configure inline tap monitor-only mode for some contexts, and regular inline mode for others.

The following figure shows the traffic flow when operating in inline tap mode.

ASA FirePOWER Management Access

There are two separate layers of access for managing an ASA FirePOWER module: initial configuration (and subsequent troubleshooting) and policy management.

- Initial Configuration, page 26-4
- Policy Configuration and Management, page 26-5

Initial Configuration

For initial configuration, you must use the CLI on the ASA FirePOWER module. For information on the default management addresses, see Default Settings, page 26-7.

To access the CLI, you can use the following methods:

- ASA 5585-X:

    - ASA FirePOWER console port—The console port on the module is a separate external console port.

    - ASA FirePOWER Management 1/0 interface using SSH—You can connect to the default IP address or you can use ASDM to change the management IP address and then connect using SSH. The management interface on the module is a separate external Gigabit Ethernet interface.

> **Note**   You cannot access the ASA FirePOWER hardware module CLI over the ASA backplane using the **session** command.

- ASA 5512-X through ASA 5555-X:

    - ASA session over the backplane—If you have CLI access to the ASA, then you can session to the module and access the module CLI.

    - ASA FirePOWER Management 0/0 interface using SSH—You can connect to the default IP address or you can use ASDM to change the management IP address and then connect using SSH. These models run the ASA FirePOWER module as a software module. The ASA FirePOWER management interface shares the Management 0/0 interface with the ASA. Separate MAC addresses and IP addresses are supported for the ASA and ASA FirePOWER

module. You must perform configuration of the ASA FirePOWER IP address within the ASA FirePOWER operating system (using the CLI or ASDM). However, physical characteristics (such as enabling the interface) are configured on the ASA. You can remove the ASA interface configuration (specifically the interface name) to dedicate this interface as an ASA FirePOWER-only interface. This interface is management-only.

Policy Configuration and Management

After you perform initial configuration, configure the ASA FirePOWER security policy using FireSIGHT Management Center. Then configure the ASA policy for sending traffic to the ASA FirePOWER module using ASDM or Cisco Security Manager.

Compatibility with ASA Features

The ASA includes many advanced application inspection features, including HTTP inspection. However, the ASA FirePOWER module provides more advanced HTTP inspection than the ASA provides, as well as additional features for other applications, including monitoring and controlling application usage.

To take full advantage of the ASA FirePOWER module features, see the following guidelines for traffic that you send to the ASA FirePOWER module:

- Do not configure ASA inspection on HTTP traffic.

- Do not configure Cloud Web Security (ScanSafe) inspection. If you configure both ASA FirePOWER inspection and Cloud Web Security inspection for the same traffic, the ASA only performs ASA FirePOWER inspection.

- Other application inspections on the ASA are compatible with the ASA FirePOWER module, including the default inspections.

- Do not enable the Mobile User Security (MUS) server; it is not compatible with the ASA FirePOWER module.

- If you enable failover, when the ASA fails over, any existing ASA FirePOWER flows are transferred to the new ASA. The ASA FirePOWER module in the new ASA begins inspecting the traffic from that point forward; old inspection states are not transferred.

Licensing Requirements for the ASA FirePOWER Module

Model	License Requirement
ASAv	Standard or Premium License.
All other models	Base License.

The ASA FirePOWER module and FireSIGHT Management Center require additional licenses. See the Licensing chapter of the *FireSIGHT System User Guide* or the online help in FireSIGHT Management Center for more information.

Guidelines and Limitations

Context Mode Guidelines

Supported in multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Failover Guidelines

Does not support failover directly; when the ASA fails over, any existing ASA FirePOWER flows are transferred to the new ASA. The ASA FirePOWER module in the new ASA begins inspecting the traffic from that point forward; old inspection states are not transferred.

You are responsible for maintaining consistent policies on the ASA FirePOWER modules in the high-availability ASA pair (using FireSIGHT Management Center) to ensure consistent failover behavior.

ASA Clustering Guidelines

Does not support clustering directly, but you can use these modules in a cluster. You are responsible for maintaining consistent policies on the ASA FirePOWER modules in the cluster using FireSIGHT Management Center. Do not use different ASA-interface-based zone definitions for devices in the cluster.

IPv6 Guidelines

Supports IPv6.

Model Guidelines

- Supported on the ASA 5585-X (as a hardware module) and 5512-X through ASA 5555-X (as a software module). See the *Cisco ASA Compatibility Matrix* for more information:

  http://www.cisco.com/en/US/docs/security/asa/compatibility/asamatrx.html

- For the 5512-X through ASA 5555-X, you must install a Cisco solid state drive (SSD). For more information, see the ASA 5500-X hardware guide.

Additional Guidelines and Limitations

- See Compatibility with ASA Features, page 26-5.

- You cannot change the software type installed on the hardware module; if you purchase an ASA FirePOWER module, you cannot later install other software on it.

- You cannot configure both normal inline mode and inline tap monitor-only mode at the same time on the ASA. Only one type of security policy is allowed. In multiple context mode, you cannot configure inline tap monitor-only mode for some contexts, and regular inline mode for others.

Default Settings

The following table lists the default settings for the ASA FirePOWER module.

Table 26-1   ASA FirePOWER Default Network Parameters

Parameters	Default
Management IP address	• System software image: 192.168.45.45/24 • Boot image: – ASA 5585-X: Management 1/0 192.168.8.8/24 – ASA 5512-X through ASA 5555-X: Management 0/0 192.168.1.2/24
Gateway	• System software image: none • Boot image: – ASA 5585-X: 192.168.8.1/24 – ASA 5512-X through ASA 5555-X: 192.168.1.1/24
SSH or session Username	admin
Password	• System software image: **Sourcefire** • Boot image: **Admin123**

Configuring the ASA FirePOWER Module

This section describes how to configure the ASA FirePOWER module.

- Task Flow for the ASA FirePOWER Module, page 26-8
- Connecting the ASA FirePOWER Management Interface, page 26-9
- (ASA 5512-X through 5555-X) Installing or Reimaging the Software Module, page 26-11
- Changing the ASA FirePOWER Management IP Address, page 26-15
- Configuring Basic ASA FirePOWER Settings at the ASA FirePOWER CLI, page 26-16
- Adding ASA FirePOWER to the FireSIGHT Management Center, page 26-17
- Configuring the Security Policy on the ASA FirePOWER Module, page 26-18
- Redirecting Traffic to the ASA FirePOWER Module, page 26-19

Task Flow for the ASA FirePOWER Module

Configuring the ASA FirePOWER module is a process that includes configuration of the ASA FirePOWER security policy on the ASA FirePOWER module and then configuration of the ASA to send traffic to the ASA FirePOWER module. To configure the ASA FirePOWER module, perform the following steps:

Step 1   Cable the ASA FirePOWER management interfaces and optionally, the console interface. See Connecting the ASA FirePOWER Management Interface, page 26-9.

Step 2   (ASA 5512-X through ASA 5555-X) Install the software module. See (ASA 5512-X through 5555-X) Installing or Reimaging the Software Module, page 26-11.

Step 3   (ASA 5585-X) Configure the ASA FirePOWER module management IP address for initial SSH access. See Changing the ASA FirePOWER Management IP Address, page 26-15.

Step 4   On the ASA FirePOWER module, configure basic settings. See Configuring Basic ASA FirePOWER Settings at the ASA FirePOWER CLI, page 26-16.

Step 5   Identify the FireSIGHT Management Center that will manage the device. See Adding ASA FirePOWER to the FireSIGHT Management Center, page 26-17.

Step 6   On the ASA FirePOWER module, configure the security policy using FireSIGHT Management Center. See Configuring the Security Policy on the ASA FirePOWER Module, page 26-18.

Step 7   On the ASA, identify traffic to divert to the ASA FirePOWER module. See Redirecting Traffic to the ASA FirePOWER Module, page 26-19.

Connecting the ASA FirePOWER Management Interface

In addition to providing management access to the ASA FirePOWER module, the ASA FirePOWER management interface needs access to an HTTP proxy server or a DNS server and the Internet for signature updates and more. This section describes recommended network configurations. Your network may differ.

ASA 5585-X (Hardware Module)

The ASA FirePOWER module includes a separate management and console interface from the ASA. For initial setup, you can connect with SSH to the ASA FirePOWER Management 1/0 interface using the default IP address. If you cannot use the default IP address, you can either use the console port or use ASDM to change the management IP address so you can use SSH. (See Changing the ASA FirePOWER Management IP Address, page 26-15.)

If you have an inside router

If you have an inside router, you can route between the management network, which can include both the ASA Management 0/0 and ASA FirePOWER Management 1/0 interfaces, and the ASA inside network for Internet access. Be sure to also add a route on the ASA to reach the Management network through the inside router.
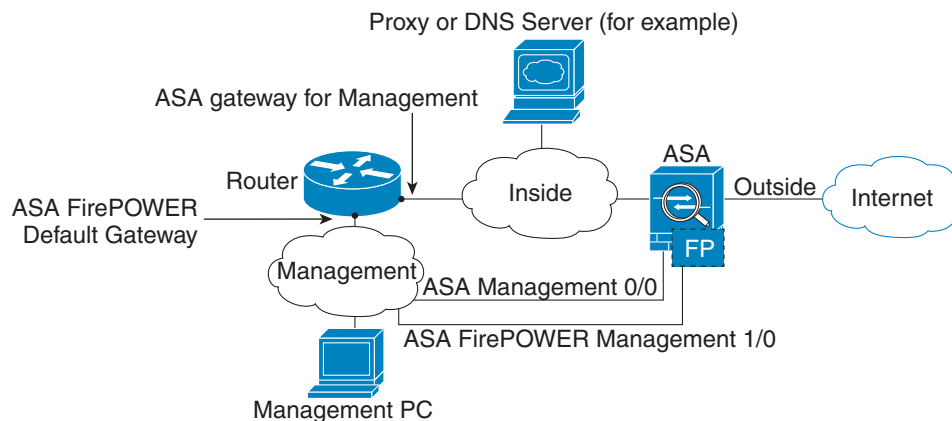
If you do not have an inside router

If you have only one inside network, then you cannot also have a separate management network, which would require an inside router to route between the networks. In this case, you can manage the ASA from the inside interface instead of the Management 0/0 interface. Because the ASA FirePOWER module is a separate device from the ASA, you can configure the ASA FirePOWER Management 1/0 address to be on the same network as the inside interface.

ASA 5512-X through ASA 5555-X (Software Module)

These models run the ASA FirePOWER module as a software module, and the ASA FirePOWER management interface shares the Management 0/0 interface with the ASA. For initial setup, you can connect with SSH to the ASA FirePOWER default IP address. If you cannot use the default IP address, you can either session to the ASA FirePOWER over the backplane or use ASDM to change the management IP address so you can use SSH.

If you have an inside router

If you have an inside router, you can route between the Management 0/0 network, which includes both the ASA and ASA FirePOWER management IP addresses, and the inside network for Internet access. Be sure to also add a route on the ASA to reach the Management network through the inside router.

If you do not have an inside router

If you have only one inside network, then you cannot also have a separate management network. In this case, you can manage the ASA from the inside interface instead of the Management 0/0 interface. If you remove the ASA-configured name from the Management 0/0 interface, you can still configure the ASA FirePOWER IP address for that interface. Because the ASA FirePOWER module is essentially a separate device from the ASA, you *can* configure the ASA FirePOWER management address to be on the same network as the inside interface.

Note   You must remove the ASA-configured name for Management 0/0; if it is configured on the ASA, then the ASA FirePOWER address must be on the same network as the ASA, and that excludes any networks already configured on other ASA interfaces. If the name is not configured, then the ASA FirePOWER address can be on any network, for example, the ASA inside network.

(ASA 5512-X through 5555-X) Installing or Reimaging the Software Module

If you purchase the ASA with the ASA FirePOWER module, the module software and required solid state drives (SSDs) come pre-installed and ready to configure. If you want to add the ASA FirePOWER software module to an existing ASA, or need to replace the SSD, you need to install the ASA FirePOWER boot software, partition the SSD, and install the system software according to this procedure.

Reimaging the module is the same procedure, except you should first uninstall the ASA FirePOWER module. You would reimage a system if you replace an SSD.

For information on how to physically install the SSD, see the ASA hardware guide.

Prerequisites

- The free space on flash (disk0) should be at least 3GB plus the size of the boot software.

- In multiple context mode, perform this procedure in the system execution space.

- You must shut down any other software module that you might be running; the device can run a single software module at a time. You must do this from the ASA CLI. For example, the following commands shut down and uninstall the IPS software module, and then reload the ASA; the commands to remove the CX module are the same, except use the **cxsc** keyword instead of **ips**.

```
hostname# sw-module module ips shutdown
hostname# sw-module module ips uninstall
hostname# reload
```

> **Note**   If you have an active service policy redirecting traffic to an IPS or CX module, you must remove that policy. For example, if the policy is a global one, you would use **no service-policy** *ips_policy* **global**. You can remove the policies using CLI or ASDM.

- When reimaging the module, use the same shutdown and uninstall commands to remove the old image. For example, **sw-module module sfr uninstall**.

- Obtain both the ASA FirePOWER Boot Image and System Software packages from Cisco.com.

Detailed Steps

Step 1   Download the boot image to the device. Do not transfer the system software; it is downloaded later to the SSD. You have the following options:

- ASDM—First, download the boot image to your workstation, or place it on an FTP, TFTP, HTTP, HTTPS, SMB, or SCP server. Then, in ASDM, choose **Tools > File Management**, and then choose the appropriate **File Transfer** command, either **Between Local PC and Flash** or **Between Remote Server and Flash**. Transfer the boot software to disk0 on the ASA.

- ASA CLI—First, place the boot image on a TFTP, FTP, HTTP, or HTTPS server, then use the **copy** command to download it to flash. The following example uses TFTP; replace <TFTP Server> with your server's IP address or host name.

```
ciscoasa# copy tftp://<TFTP SERVER>/asasfr-5500x-boot-5.3.1-58.img
disk0:/asasfr-5500x-boot-5.3.1-58.img
```

Step 2   Download the ASA FirePOWER system software from Cisco.com to an HTTP, HTTPS, or FTP server accessible from the ASA FirePOWER management interface.

Step 3   Set the ASA FirePOWER module boot image location in ASA disk0 by entering the following command:

```
hostname# sw-module module sfr recover configure image disk0:file_path
```

> **Note** If you get a message like "ERROR: Another service (cxsc) is running, only one service is allowed to run at any time," it means that you already have a different software module configured. You must shut it down and remove it to install a new module as described in the prerequisites section above.

Example:

```
hostname# sw-module module sfr recover configure image
disk0:asasfr-5500x-boot-5.3.1-58.img
```

Step 4   Load the ASA FirePOWER boot image by entering the following command:

```
hostname# sw-module module sfr recover boot
```

Step 5   Wait approximately 5-15 minutes for the ASA FirePOWER module to boot up, and then open a console session to the now-running ASA FirePOWER boot image. You might need to press enter after opening the session to get to the login prompt. The default username is **admin** and the default password is **Admin123**.

```
hostname# session sfr console
Opening console session with module sfr.
Connected to module sfr. Escape character sequence is 'CTRL-^X'.

Cisco ASA SFR Boot Image 5.3.1
asasfr login: admin
Password: Admin123
```

> **Tip**   If the module boot has not competed, the **session** command will fail with a message about not being able to connect over ttyS1. Wait and try again.

Step 6   Use the **setup** command to configure the system so that you can install the system software package.

```
asasfr-boot> setup

                    Welcome to SFR Setup
                     [hit Ctrl-C to abort]
                Default values are inside []
```

You are prompted for the following. Note that the management address and gateway, and DNS information, are the key settings to configure.

- Host name—Up to 65 alphanumeric characters, no spaces. Hyphens are allowed.

- Network address—You can set static IPv4 or IPv6 addresses, or use DHCP (for IPv4) or IPv6 stateless autoconfiguration.

- DNS information—You must identify at least one DNS server, and you can also set the domain name and search domain.

- NTP information—You can enable NTP and configure the NTP servers, for setting system time.

Step 7   Install the System Software image using the **system install** command:

system install [**noconfirm**] *url*

Include the **noconfirm** option if you do not want to respond to confirmation messages. Use an HTTP, HTTPS, or FTP URL; if a username and password are required, you will be prompted to supply them.

When installation is complete, the system reboots. Allow 10 or more minutes for application component installation and for the ASA FirePOWER services to start. (The **show module sfr** output should show all processes as Up.)

For example:

```
asasfr-boot> system install http://asasfr-sys-5.3.1-44.pkg
Verifying
Downloading
Extracting
Package Detail
        Description:                    Cisco ASA-FirePOWER 5.3.1-44 System Install
        Requires reboot:               Yes

Do you want to continue with upgrade? [y]: y
Warning: Please do not interrupt the process or turn off the system.
Doing so might leave system in unusable state.

Upgrading
Starting upgrade process ...
Populating new system image

Reboot is required to complete the upgrade. Press 'Enter' to reboot the system.
(press Enter)
Broadcast message from root (ttyS1) (Mon Feb 17 19:28:38 2014):

The system is going down for reboot NOW!
Console session with module sfr terminated.
```

Step 8   Open a session to the ASA FirePOWER module. You will see a different login prompt because you are logging into the fully functional module.

```
asa3# session sfr
Opening command session with module sfr.
Connected to module sfr. Escape character sequence is 'CTRL-^X'.

Sourcefire ASA5555 v5.3.1 (build 44)
Sourcefire3D login:
```

Step 9   Log in with the username **admin** and the password **Sourcefire**.

Step 10   Complete the system configuration as prompted.

You must first read and accept the end user license agreement (EULA). Then change the admin password, then configure the management address and DNS settings, as prompted. You can configure both IPv4 and IPv6 management addresses. For example:

```
System initialization in progress.  Please stand by.
You must change the password for 'admin' to continue.
Enter new password: <new password>
Confirm new password: <repeat password>
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]: 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]: 10.86.118.3
Enter an IPv4 netmask for the management interface [255.255.255.0]: 255.255.252.0
Enter the IPv4 default gateway for the management interface []: 10.86.116.1
Enter a fully qualified hostname for this system [Sourcefire3D]: asasfr.example.com
Enter a comma-separated list of DNS servers or 'none' []: 10.100.10.15,
10.120.10.14
Enter a comma-separated list of search domains or 'none' [example.net]: example.com
If your networking information has changed, you will need to reconnect.
```

```
For HTTP Proxy configuration, run 'configure network http-proxy'
(Wait for the system to reconfigure itself.)

This sensor must be managed by a Defense Center.  A unique alphanumeric
registration key is always required.  In most cases, to register a sensor
to a Defense Center, you must provide the hostname or the IP address along
with the registration key.
'configure manager add [hostname | ip address ] [registration key ]'

However, if the sensor and the Defense Center are separated by a NAT device,
you must enter a unique NAT ID, along with the unique registration key.
'configure manager add DONTRESOLVE [registration key ] [ NAT ID ]'

Later, using the web interface on the Defense Center, you must use the same
registration key and, if necessary, the same NAT ID when you add this
sensor to the Defense Center.
```

Step 11 Identify the FireSIGHT Management Center appliance that will manage this device using the **configure manager add** command.

You come up with a registration key, which you will then use in FireSIGHT Management Center when you add the device to its inventory. The following example shows the simple case. When there is a NAT boundary, the command is different; see Adding ASA FirePOWER to the FireSIGHT Management Center, page 26-17.

```
> configure manager add 10.89.133.202 123456
Manager successfully configured.
```

Step 12 Log into the FireSIGHT Management Center using an HTTPS connection in a browser, using the hostname or address entered above. For example, https://DC.example.com.

Use the Device Management (**Devices > Device Management**) page to add the device. For more information, see the online help or the Managing Devices chapter in the *FireSIGHT System User Guide*.

Tip You also configure NTP and time settings through FireSIGHT Management Center. Use the Time Synchronization settings when editing the local policy from the **System > Local > System Policy** page.

Changing the ASA FirePOWER Management IP Address

If you cannot use the default management IP address, then you can set the management IP address from the ASA. After you set the management IP address, you can access the ASA FirePOWER module using SSH to perform additional setup.

If you already configured the management address during initial system setup through the ASA FirePOWER CLI, as described in Configuring Basic ASA FirePOWER Settings at the ASA FirePOWER CLI, page 26-16, then it is not necessary to configure it through the ASA CLI or ASDM.

Note For a software module, you can access the ASA FirePOWER CLI to perform setup by sessioning from the ASA CLI; you can then set the ASA FirePOWER management IP address as part of setup. For a hardware module, you can complete the initial setup through the Console port.

Guidelines

In multiple context mode, perform this procedure in the system execution space.

Detailed Steps

Step 1   In ASDM, choose **Wizards > Startup Wizard**.

Step 2   Click **Next** to advance through the initial screens until you reach the ASA FirePOWER Basic Configuration screen.

Step 3   Enter the new management IP address, subnet mask, and default gateway.

You must also accept the end user license agreement.

Step 4   Click **Finish** to skip the remaining screens, or click **Next** to advance through the remaining screens and complete the wizard.

Configuring Basic ASA FirePOWER Settings at the ASA FirePOWER CLI

You must configure basic network settings and other parameters on the ASA FirePOWER module before you can configure your security policy. This procedure assumes you have the full system software installed (not just the boot image), either after you installed it directly, or because it is already installed on a hardware module.

Tip   This procedure also assumes that you are performing an initial configuration. During initial configuration, you are prompted for these settings. If you need to change these settings later, use the various **configure network** commands to change the individual settings. For more information on the **configure network** commands, use the **?** command for help, and see the *FireSIGHT System User Guide*, or the online help in FireSIGHT Management Center.

Detailed Steps

Step 1   Do one of the following:

- (All models) Use SSH to connect to the ASA FirePOWER management IP address.

- (ASA 5512-X through ASA 5555-X) Open a session to the module from the ASA CLI (see the "Getting Started" chapter in the general operations configuration guide to access the ASA CLI). In multiple context mode, session from the system execution space.

      hostname# **session sfr**

Step 2   Log in with the username **admin** and the password **Sourcefire**.

Step 3   Complete the system configuration as prompted.

You must first read and accept the end user license agreement (EULA). Then change the admin password, then configure the management address and DNS settings, as prompted. You can configure both IPv4 and IPv6 management addresses. The configuration is complete when you see the message that says the sensor must be managed by a FireSIGHT Management Center.

For example:

```
System initialization in progress.  Please stand by.
You must change the password for 'admin' to continue.
Enter new password: <new password>
Confirm new password: <repeat password>
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]: 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]: 10.86.118.3
Enter an IPv4 netmask for the management interface [255.255.255.0]: 255.255.252.0
Enter the IPv4 default gateway for the management interface []: 10.86.116.1
Enter a fully qualified hostname for this system [Sourcefire3D]: asasfr.example.com
Enter a comma-separated list of DNS servers or 'none' []: 10.100.10.15,
10.120.10.14
Enter a comma-separated list of search domains or 'none' [example.net]: example.com
If your networking information has changed, you will need to reconnect.
For HTTP Proxy configuration, run 'configure network http-proxy'
(Wait for the system to reconfigure itself.)

This sensor must be managed by a Defense Center.  A unique alphanumeric
registration key is always required.  In most cases, to register a sensor
to a Defense Center, you must provide the hostname or the IP address along
with the registration key.
'configure manager add [hostname | ip address ] [registration key ]'

However, if the sensor and the Defense Center are separated by a NAT device,
you must enter a unique NAT ID, along with the unique registration key.
'configure manager add DONTRESOLVE [registration key ] [ NAT ID ]'

Later, using the web interface on the Defense Center, you must use the same
registration key and, if necessary, the same NAT ID when you add this
sensor to the Defense Center.
```

Step 4   Now you must identify the FireSIGHT Management Center that will manage this device, as explained in Adding ASA FirePOWER to the FireSIGHT Management Center, page 26-17.

Adding ASA FirePOWER to the FireSIGHT Management Center

You must register the ASA FirePOWER module to a FireSIGHT Management Center, which is the application you use to configure the policies on the module. FireSIGHT Management Center is also known as Defense Center.

To register a device, use the **configure manager add** command. A unique alphanumeric registration key is always required to register a device to a FireSIGHT Management Center. This is a simple key that you specify, and is not the same as a license key.

In most cases, you must provide the FireSIGHT Management Center's hostname or the IP address along with the registration key, for example:

```
configure manager add DC.example.com my_reg_key
```

However, if the device and the FireSIGHT Management Center are separated by a NAT device, enter a unique NAT ID along with the registration key, and specify DONTRESOLVE instead of the hostname, for example:

```
configure manager add DONTRESOLVE my_reg_key my_nat_id
```

Detailed Steps

Step 1   Do one of the following:

- (All models) Use SSH to connect to the ASA FirePOWER management IP address.

- (ASA 5512-X through ASA 5555-X) Open a session to the module from the ASA CLI (see the "Getting Started" chapter in the general operations configuration guide to access the ASA CLI). In multiple context mode, session from the system execution space.

  ```
  hostname# session sfr
  ```

Step 2   Log in with the username **admin** or another username that has the CLI configuration (Administrator) access level.

Step 3   At the prompt, register the device to a FireSIGHT Management Center using the **configure manager add** command, which has the following syntax:

configure manager add {*hostname* | *IPv4_address* | *IPv6_address* | **DONTRESOLVE**} *reg_key* [*nat_id*]

where:

- {*hostname* | *IPv4_address* | *IPv6_address* | **DONTRESOLVE**} specifies either the fully qualified host name or IP address of the FireSIGHT Management Center. If the FireSIGHT Management Center is not directly addressable, use DONTRESOLVE.

- *reg_key* is the unique alphanumeric registration key required to register a device to the FireSIGHT Management Center.

- *nat_id* is an optional alphanumeric string used during the registration process between the FireSIGHT Management Center and the device. It is required if the hostname is set to DONTRESOLVE.

Step 4   Log into the FireSIGHT Management Center using an HTTPS connection in a browser, using the hostname or address entered above. For example, https://DC.example.com.

Use the Device Management (**Devices > Device Management**) page to add the device. For more information, see the online help or the Managing Devices chapter in the *FireSIGHT System User Guide*.

Configuring the Security Policy on the ASA FirePOWER Module

You use FireSIGHT Management Center to configure the security policy on the ASA FirePOWER module. The security policy controls the services provided by the module, such as Next Generation IPS filtering and application filtering. You cannot configure the policy through the ASA FirePOWER CLI, the ASA CLI, or ASDM.

To open FireSIGHT Management Center, use a web browser to open the following URL:

https://*DC_address*

Where *DC_address* is the DNS name or IP address of the manager you defined in Adding ASA FirePOWER to the FireSIGHT Management Center, page 26-17. For example, https://dc.example.com.

For information about how to configure the security policy, see the *FireSIGHT System User Guide* or the online help in FireSIGHT Management Center.

Tip   You can also open FireSIGHT Management Center from the ASA FirePOWER Status dashboard in ASDM. Choose **Home > ASA FirePOWER Status**, and click the link at the bottom of the dashboard.

Redirecting Traffic to the ASA FirePOWER Module

Redirect traffic to the ASA FirePOWER module by creating a service policy that identifies specific traffic that you want to send. ASA policies, such as access rules, are applied to the traffic before it is redirected to the module.

You can configure your device in either an inline or inline tap monitor-only deployment.

- In an inline deployment, the actual traffic is sent to the device, and the device's policy affects what happens to the traffic. After dropping undesired traffic and taking any other actions applied by policy, the traffic is returned to the ASA for further processing and ultimate transmission.

- In an inline tap deployment, a copy of the traffic is sent to the device, but it is not returned to the ASA. Inline tap mode lets you see what the device would have done to traffic, and lets you evaluate the content of the traffic, without impacting the network.

Note   You cannot configure both monitor-only mode and normal inline mode at the same time on the ASA. Only one type of security policy is allowed. In multiple context mode, you cannot configure monitor-only mode for some contexts, and regular inline mode for others.

Prerequisites

- If you have an active service policy redirecting traffic to an IPS or CX module (that you replaced with the ASA FirePOWER), you must remove that policy before you configure the ASA FirePOWER service policy.

- Be sure to configure consistent policies on the ASA and the ASA FirePOWER. Both policies should reflect the inline or inline tap mode of the traffic.

- In multiple context mode, perform this procedure within each security context.

Detailed Steps

Step 1   Choose **Configuration > Firewall > Service Policy Rules**.

Step 2   Choose **Add > Add Service Policy Rule**. The Add Service Policy Rule Wizard - Service Policy dialog box appears.

Step 3   Complete the Service Policy dialog box as desired. See the ASDM online help for more information about these screens.

Step 4   Click **Next**. The Add Service Policy Rule Wizard - Traffic Classification Criteria dialog box appears.

Step 5   Complete the Traffic Classification Criteria dialog box as desired. See the ASDM online help for more information about these screens.

Step 6   Click **Next** to show the Add Service Policy Rule Wizard - Rule Actions dialog box.

Step 7   Click the **ASA FirePOWER Inspection** tab.

Step 8   Check the **Enable ASA FirePOWER for this traffic flow** check box.

Step 9   In the If ASA FirePOWER Card Fails area, click one of the following:

- **Permit traffic**—Sets the ASA to allow all traffic through, uninspected, if the module is unavailable.

- **Close traffic**—Sets the ASA to block all traffic if the module is unavailable.

Step 10   (Optional) Check **Monitor-only** to send a read-only copy of traffic to the module, i.e. inline tap mode.

By default, the traffic is sent in inline mode. Be sure to configure consistent policies on the ASA and the ASA FirePOWER. Both policies should reflect the inline or monitor-only of the traffic.

Step 11   Click **Finish** and then **Apply**.

Repeat this procedure to configure additional traffic flows as desired.

Managing the ASA FirePOWER Module

This section includes procedures that help you manage the module.

Resetting the Password

If you forget the password for admin user, another user with CLI Configuration permissions can log in and change the password.

If there are no other users with the required permissions, you can reset the admin password from the ASA using the **session do** command.

Tip The password-reset option on the ASA hw-module and sw-module commands does not work with ASA FirePOWER.

Guidelines

In multiple context mode, perform this procedure in the system execution space.

Detailed Steps

Command	Purpose
For a hardware module (ASA 5585-X): **session 1 do password-reset** For a software module (ASA 5512-X through ASA 5555-X): **session sfr do password-reset** **Example:** hostname# session sfr do password-reset	Resets the module password for the user **admin** to the default, **Sourcefire**.

Reloading or Resetting the Module

To reload or reset the module, enter one of the following commands at the ASA CLI.

Guidelines

In multiple context mode, perform this procedure in the system execution space.

Detailed Steps

Command	Purpose

For a hardware module (ASA 5585-X):
 	
`hw-module module 1 reload`
 	
For a software module (ASA 5512-X through ASA 5555-X):
 	
`sw-module module sfr reload`
 	
Example:
hostname# hw-module module 1 reload
	Reloads the module software.

For a hardware module:
 	
`hw-module module 1 reset`
 	
For a software module:
 	
`sw-module module sfr reset`
 	
Example:
hostname# hw-module module 1 reset
	Performs a reset, and then reloads the module.

Shutting Down the Module

Shutting down the module software prepares the module to be safely powered off without losing configuration data. To gracefully shut down the module, perform the following steps at the ASA CLI.

Note   If you reload the ASA, the module is not automatically shut down, so we recommend shutting down the module before reloading the ASA.

Guidelines

In multiple context mode, perform this procedure in the system execution space.

Detailed Steps

Command	Purpose

For a hardware module (ASA 5585-X):
 	
`hw-module module 1 shutdown`
 	
For a software module (ASA 5512-X through ASA 5555-X):
 	
`sw-module module sfr shutdown`
 	
Example:
`hostname# hw-module module 1 shutdown`
	Shuts down the module.

(ASA 5512-X through ASA 5555-X) Uninstalling a Software Module Image

To uninstall a software module image and associated configuration, perform the following steps.

Guidelines

In multiple context mode, perform this procedure in the system execution space.

Detailed Steps

	Command	Purpose

Step 1	`sw-module module sfr uninstall`
	
	Example:
	`hostname# sw-module module sfr uninstall`
	`Module sfr will be uninstalled. This will`
	`completely remove the`
	`disk image associated with the sw-module`
	`including any configuration`
	`that existed within it.`
	
	`Uninstall module <id>? [confirm]`
		Permanently uninstalls the software module image and associated configuration.

Step 2	`reload`
	
	Example:
	`hostname# reload`
		Reloads the ASA. You must reload the ASA before you can install a new module type.

(ASA 5512-X through ASA 5555-X) Sessioning to the Module From the ASA

To access the ASA FirePOWER software module CLI from the ASA, you can session from the ASA. You can either session to the module (using Telnet) or create a virtual console session. A console session might be useful if the control plane is down and you cannot establish a Telnet session.

Use the ASA FirePOWER CLI to configure basic network settings and to troubleshoot the module.

Guidelines

- In multiple context mode, perform this procedure in the system execution space.
- You can log in with any username configured on the ASA FirePOWER. Initially, the **admin** username is the only one configured (and it is always available). The initial default username is **Sourcefire** for the full image, and **Admin123** for the boot image.

Detailed Steps

Command	Purpose
Telnet session. **session sfr** **Example:** hostname# session sfr Opening command session with module sfr. Connected to module sfr. Escape character sequence is 'CTRL-^X'. asasfr login: admin Password:	Accesses the module using Telnet. You are prompted for the username and password. When in the ASA FirePOWER CLI, to exit back to the ASA CLI: • Enter any command that would log you out of the module, such as **logout** or **exit**. • Press **Ctrl-Shift-6, x**.
Console session. **session sfr console** **Example:** hostname# session sfr console Opening console session with module sfr. Connected to module sfr. Escape character sequence is 'CTRL-^X'. asasfr login: admin Password:	Accesses the module console. You are prompted for the username and password. The only way out of a console session is to press **Ctrl-Shift-6, x**. Logging out of the module leaves you at the module login prompt. **Note** Do not use this command in conjunction with a terminal server where **Ctrl-Shift-6, x** is the escape sequence to return to the terminal server prompt. **Ctrl-Shift-6, x** is also the sequence to escape the ASA FirePOWER console and return to the ASA prompt. Therefore, if you try to exit the ASA FirePOWER console in this situation, you instead exit all the way to the terminal server prompt. If you reconnect the terminal server to the ASA, the ASA FirePOWER console session is still active; you can never exit to the ASA prompt. You must use a direct serial connection to return the console to the ASA prompt. Use the **session sfr** command instead.

Reimaging the 5585-X ASA FirePOWER Hardware Module

If you need to reimage the ASA FirePOWER hardware module in an ASA 5585-X appliance for any reason, you need to install both the Boot Image and a System Software package, in that order. You must install both packages to have a functioning system. Under normal circumstances, you do not need to reimage the system to install upgrade packages.

To install the boot image, you need to TFTP boot the image from the Management-0 port on the ASA FirePOWER SSP by logging into the module's Console port. Because the Management-0 port is on an SSP in the first slot, it is also known as Management1/0, but rommon recognizes it as Management-0 or Management0/1.

To accomplish a TFTP boot, you must:

- Place the software image on a TFTP server that can be accessed through the Management1/0 interface on the ASA FirePOWER.
- Connect Management1/0 to the network. You must use this interface to TFTP boot the Boot Image.
- Configure rommon variables. Press Esc to interrupt the auto-boot process so that you can configure rommon variables.

Once the boot image is installed, you install the System Software package. You must place the package on an HTTP, HTTPS, or FTP server that is accessible from the ASA FirePOWER.

The following procedure explains how to install the boot image and then install the System Software package.

Detailed Steps

Step 1   Connect to the Console port. Use the console cable included with the ASA product 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 ASA for more information about the console cable.

Step 2   Enter the **system reboot** command to reload the system.

Step 3   When prompted, break out of the boot by pressing Esc. If you see grub start to boot the system, you have waited too long.

This will place you at the rommon prompt.

Step 4   At the rommon prompt, enter **set** and configure the following parameters:

- ADDRESS—The management IP address of the module.
- SERVER—The IP address of the TFTP server.
- GATEWAY—The gateway address to the TFTP server. If the TFTP server is directly attached to Management1/0, use the IP address of the TFTP server. If the TFTP server and management address are on the same subnet, do not configure the gateway or TFTP boot will fail.
- IMAGE—The Boot Image path and image name on the TFTP server. For example, if you place the file on the TFTP server in /tftpboot/images/filename.img, the IMAGE value is images/filename.img.

For example:

```
ADDRESS=10.5.190.199
SERVER=10.5.11.170
GATEWAY=10.5.1.1
IMAGE=asasfr-boot-5.3.1-26-54.img
```

Step 5   Enter **sync** to save the settings.

Step 6   Enter **tftp** to initiate the download and boot process.

You will see ! marks to indicate progress. When the boot completes after several minutes, you will see a login prompt.

Step 7   Log in as **admin**, with the password **Admin123**.

Step 8   Use the **setup** command to configure the system so that you can install the system software package.

You are prompted for the following. Note that the management address and gateway, and DNS information, are the key settings to configure.

- Host name—Up to 65 alphanumeric characters, no spaces. Hyphens are allowed.
- Network address—You can set static IPv4 or IPv6 addresses, or use DHCP (for IPv4) or IPv6 stateless autoconfiguration.
- DNS information—You must identify at least one DNS server, and you can also set the domain name and search domain.
- NTP information—You can enable NTP and configure the NTP servers, for setting system time.

Step 9   Install the System Software image using the **system install** command:

system install [**noconfirm**] *url*

Include the **noconfirm** option if you do not want to respond to confirmation messages.

When installation is complete, the system reboots. Allow 10 or more minutes for application component installation and for the ASA FirePOWER services to start.

For example:

```
asasfr-boot> system install http://asasfr-sys-5.3.1-54.pkg
```

Step 10   When the boot completes, log in as **admin** with the password **Sourcefire**.

Complete the system configuration as prompted.

You must first read and accept the end user license agreement (EULA). Then change the admin password, then configure the management address and DNS settings, as prompted. You can configure both IPv4 and IPv6 management addresses.

Step 11   Identify the FireSIGHT Management Center appliance that will manage this device using the **configure manager add** command.

You come up with a registration key, which you will then use in FireSIGHT Management Center when you add the device to its inventory. The following example shows the simple case. When there is a NAT boundary, the command is different; see Adding ASA FirePOWER to the FireSIGHT Management Center, page 26-17.

```
> configure manager add 10.89.133.202 123456
Manager successfully configured.
```

Step 12   Log into the FireSIGHT Management Center using an HTTPS connection in a browser, using the hostname or address entered above. For example, https://DC.example.com.

Use the Device Management (**Devices > Device Management**) page to add the device. For more information, see the Managing Devices chapter in the *FireSIGHT System User Guide* or the online help in FireSIGHT Management Center.

Upgrading the System Software

Use FireSIGHT Management Center to apply upgrade images to the ASA FirePOWER module. Before applying an upgrade, ensure that the ASA is running the minimum required release for the new version; you might need to upgrade the ASA prior to upgrading the module.

For more information about applying upgrades, see the *FireSIGHT System User Guide* or the online help in FireSIGHT Management Center.

Monitoring the ASA FirePOWER Module

Use **Tools > Command Line Interface** to use monitoring commands.

- Showing Module Status, page 26-27
- Showing Module Statistics, page 26-27
- Monitoring Module Connections, page 26-28
- Capturing Module Traffic, page 26-30

Note   For ASA FirePOWER-related syslog messages, see the syslog messages guide. ASA FirePOWER syslog messages start with message number 434001.

Showing Module Status

From the Home page, you can select the ASA FirePOWER Status tab to view information about the module. This includes module information, such as the model, serial number, and software version, and module status, such as the application name and status, data plane status, and overall status. If the module is registered to a FireSIGHT Management Center, you can click the link to open the application and do further analysis and module configuration.

Showing Module Statistics

To show module statistics, enter the following command:

Command	Purpose
`show service-policy sfr`	Displays statistics and status for each service policy that includes the **sfr** command. Use **clear service-policy** to clear the counters.

Examples

The following example shows the ASA FirePOWER service policy and the current statistics as well as the module status:

```
ciscoasa# show service-policy sfr

Global policy:
  Service-policy: global_policy
    Class-map: my-sfr-class
```

```
                        SFR: card status Up, mode fail-close
                          packet input 2626422041, packet output 2626877967, drop 0, reset-drop 0, proxied 0
```

The following example shows a monitor-only policy. In this case, you should see packet input counters increasing, but the packet output counter should stay zero, because no traffic is passing back to the ASA.

```
hostname# show service-policy sfr
Global policy:
  Service-policy: global_policy
    Class-map: bypass
      SFR: card status Up, mode fail-open, monitor-only
        packet input 2626422041, packet output 0, drop 0, reset-drop 0, proxied 0
```

Monitoring Module Connections

To show connections through the ASA FirePOWER module, enter one of the following commands:

Command	Purpose
`show asp table classify domain sfr`	Shows the NP rules created to send traffic to the ASA FirePOWER module.
`show asp drop`	Shows dropped packets. The drop types are explained following this table.
`show conn`	This command already shows if a connection is being forwarded to a module by displaying the 'X - inspected by service module' flag.

Drop Reasons

The show asp drop command can include the following drop reasons related to the ASA FirePOWER module.

Frame Drops:

- sfr-bad-tlv-received—This occurs when ASA receives a packet from FirePOWER without a Policy ID TLV. This TLV must be present in non-control packets if it does not have the Standby/Active bit set in the actions field.

- sfr-request—The frame was requested to be dropped by FirePOWER due a policy on FirePOWER whereby FirePOWER would set the actions to Deny Source, Deny Destination, or Deny Pkt. If the frame should not have been dropped, review the policies on the module that are denying the flow.

- sfr-fail-close—The packet is dropped because the card is not up and the policy configured was 'fail-close' (rather than 'fail-open' which allows packets through even if the card was down). Check card status and attempt to restart services or reboot it.

- sfr-fail—The FirePOWER configuration was removed for an existing flow and we are not able to process it through FirePOWER it will be dropped. This should be very unlikely.

- sfr-malformed-packet—The packet from FirePOWER contains an invalid header. For instance, the header length may not be correct.

- sfr-ha-request—This counter is incremented when the security appliance receives a FirePOWER HA request packet, but could not process it and the packet is dropped.

- sfr-invalid-encap—This counter is incremented when the security appliance receives a FirePOWER packet with invalid message header, and the packet is dropped.

- sfr-bad-handle-received—Received Bad flow handle in a packet from FirePOWER Module, thus dropping flow. This counter is incremented, flow and packet are dropped on ASA as the handle for FirePOWER flow has changed in flow duration.

- sfr-rx-monitor-only—This counter is incremented when the security appliance receives a FirePOWER packet when in monitor-only mode, and the packet is dropped.

Flow Drops:

- sfr-request—The FirePOWER requested to terminate the flow. The actions bit 0 is set.

- reset-by-sfr—The FirePOWER requested to terminate and reset the flow. The actions bit 1 is set.

- sfr-fail-close—The flow was terminated because the card is down and the configured policy was 'fail-close'.

Examples

The following is sample output from the **show asp table classify domain sfr** command:

```
hostname# show asp table classify domain sfr

Input Table
in  id=0x7ffe60139410, priority=73, domain=sfr, deny=false
        hits=0, user_data=0x7ffe5c5932c0, cs_id=0x0, use_real_addr, flags=0x0, protocol=6
        src ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=0
        dst ip/id=0.0.0.0, mask=0.0.0.0, port=6000, tag=0, dscp=0x0
        input_ifc=outside, output_ifc=any
in  id=0x7ffe60139510, priority=73, domain=sfr, deny=false
        hits=0, user_data=0x7ffe5c5932c0, cs_id=0x0, use_real_addr, flags=0x10000,
protocol=6
        src ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=0
        dst ip/id=::/0, port=6000, tag=0, dscp=0x0
        input_ifc=outside, output_ifc=any
in  id=0x7ffe60139610, priority=73, domain=sfr, deny=false
        hits=0, user_data=0x7ffe5c5932c0, cs_id=0x0, use_real_addr, flags=0x0, protocol=6
        src ip/id=::/0, port=0, tag=0
        dst ip/id=::/0, port=6000, tag=0
        input_ifc=outside, output_ifc=any
in  id=0x7ffe5c593f20, priority=73, domain=sfr, deny=false
        hits=0, user_data=0x7ffe5c5932c0, cs_id=0x0, use_real_addr, flags=0x20000,
protocol=6
        src ip/id=::/0, port=0, tag=0
        dst ip/id=0.0.0.0, mask=0.0.0.0, port=6000, tag=0
        input_ifc=outside, output_ifc=any

Output Table:
out id=0x7ffe5c594560, priority=73, domain=sfr, deny=false
        hits=0, user_data=0x7ffe5c5932c0, cs_id=0x0, use_real_addr, flags=0x0, protocol=6
        src ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=0
        dst ip/id=0.0.0.0, mask=0.0.0.0, port=6000, tag=0, dscp=0x0
        input_ifc=any, output_ifc=outside
out id=0x7ffe5c595b80, priority=73, domain=sfr, deny=false
        hits=0, user_data=0x7ffe5c5932c0, cs_id=0x0, use_real_addr, flags=0x20000,
protocol=6
        src ip/id=::/0, port=0, tag=0
        dst ip/id=0.0.0.0, mask=0.0.0.0, port=6000, tag=0, dscp=0x0
        input_ifc=any, output_ifc=outside
out id=0x7ffe5c595400, priority=73, domain=sfr, deny=false
        hits=0, user_data=0x7ffe5c5932c0, cs_id=0x0, use_real_addr, flags=0x10000,
protocol=6
        src ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=0
        dst ip/id=::/0, port=6000, tag=0
```

```
                    input_ifc=any, output_ifc=outside
            out id=0x7ffe5c5957c0, priority=73, domain=sfr, deny=false
                    hits=0, user_data=0x7ffe5c5932c0, cs_id=0x0, use_real_addr, flags=0x0, protocol=6
                    src ip/id=::/0, port=0, tag=0
                    dst ip/id=::/0, port=6000, tag=0
                    input_ifc=any, output_ifc=outside

            L2 - Output Table:

            L2 - Input Table:

            Last clearing of hits counters: Never
```

Capturing Module Traffic

To configure and view packet captures for the module, enter one of the following commands:

Command	Purpose
`capture` *name* `interface asa_dataplane`	Captures packets between module and the ASA on the backplane.
`copy capture`	Copies the capture file to a server.
`show capture`	Shows the capture at the ASA console.

Note   Captured packets contain an additional AFBP header that your PCAP viewer might not understand; be sure to use the appropriate plug-in to view these packets.

Feature History for the ASA FirePOWER Module

The following table lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 26-2   Feature History for the ASA FirePOWER Module

Feature Name	Platform Releases	Feature Information
ASA 5585-X (all models) support for the matching ASA FirePOWER SSP hardware module. ASA 5512-X through ASA 5555-X support for the ASA FirePOWER software module.	ASA 9.2(2) ASA FirePOWER 5.3.1	The ASA FirePOWER module supplies next-generation firewall services, including Next-Generation IPS (NGIPS), Application Visibility and Control (AVC), URL filtering, and Advanced Malware Protection (AMP).You can use the module in single or multiple context mode, and in routed or transparent mode. We introduced the following screens: Home > ASA FirePOWER Status Wizards > Startup Wizard > ASA FirePOWER Basic Configuration Configuration > Firewall > Service Policy Rules > Add Service Policy Rule > Rule Actions > ASA FirePOWER Inspection

CHAPTER **27**

ASA CX Module

This chapter describes how to configure the ASA CX module that runs on the ASA.

Information About the ASA CX Module

The ASA CX module lets you enforce security based on the full context of a situation. This context includes the identity of the user (who), the application or website that the user is trying to access (what), the origin of the access attempt (where), the time of the attempted access (when), and the properties of the device used for the access (how). With the ASA CX module, you can extract the full context of a flow and enforce granular policies such as permitting access to Facebook but denying access to games on Facebook, or permitting finance employees access to a sensitive enterprise database but denying the same access to other employees.

How the ASA CX Module Works with the ASA

The ASA CX module runs a separate application from the ASA. The ASA CX module includes external management interface(s) so you can connect to the ASA CX module directly. Any data interfaces on the ASA CX module are used for ASA traffic only.

Traffic goes through the firewall checks before being forwarded to the ASA CX module. When you identify traffic for ASA CX inspection on the ASA, traffic flows through the ASA and the ASA CX module as follows:

1. Traffic enters the ASA.
2. Incoming VPN traffic is decrypted.
3. Firewall policies are applied.
4. Traffic is sent to the ASA CX module.
5. The ASA CX module applies its security policy to the traffic, and takes appropriate actions.
6. Valid traffic is sent back to the ASA; the ASA CX module might block some traffic according to its security policy, and that traffic is not passed on.
7. Outgoing VPN traffic is encrypted.
8. Traffic exits the ASA.

Figure 27-1 shows the traffic flow when using the ASA CX module. In this example, the ASA CX module automatically blocks traffic that is not allowed for a certain application. All other traffic is forwarded through the ASA.

Figure 27-1   ASA CX Module Traffic Flow in the ASA
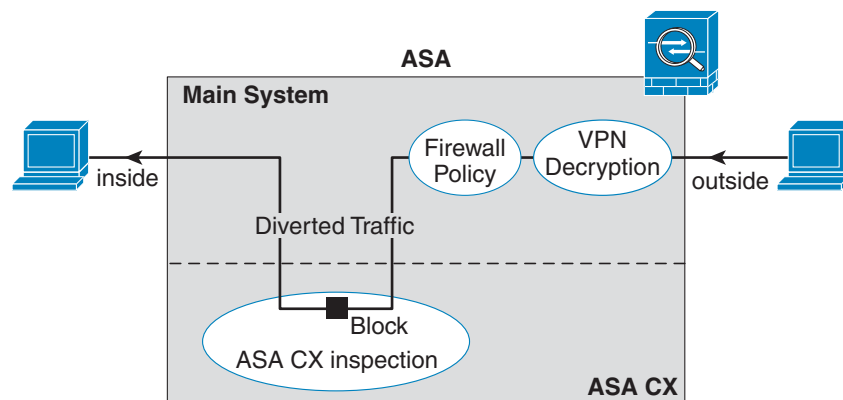
Note   If you have a connection between hosts on two ASA interfaces, and the ASA CX service policy is only configured for one of the interfaces, then all traffic between these hosts is sent to the ASA CX module, including traffic originating on the non-ASA CX interface (because the feature is bidirectional). However, the ASA only performs the authentication proxy on the interface to which the service policy is applied, because authentication proxy is applied only to ingress traffic (see Information About Authentication Proxy, page 27-5).

Monitor-Only Mode

For demonstration purposes, you can configure a service policy or a traffic-forwarding interface in monitor-only mode.

For guidelines and limitations for monitor-only mode, see Guidelines and Limitations, page 27-6.

- Service Policy in Monitor-Only Mode, page 27-3
- Traffic-Forwarding Interface in Monitor-Only Mode, page 27-3

Service Policy in Monitor-Only Mode

For testing and demonstration purposes, you can configure the ASA to send a duplicate stream of read-only traffic to the ASA CX module, so you can see how the module inspects the traffic without affecting the ASA traffic flow. In this mode, the ASA CX module inspects the traffic as usual, makes policy decisions, and generates events. However, because the packets are read-only copies, the module actions do not affect the actual traffic. Instead, the module drops the copies after inspection. Figure 27-2 shows the ASA CX module in monitor-only mode.

Figure 27-2   ASA CX Monitor-Only Mode
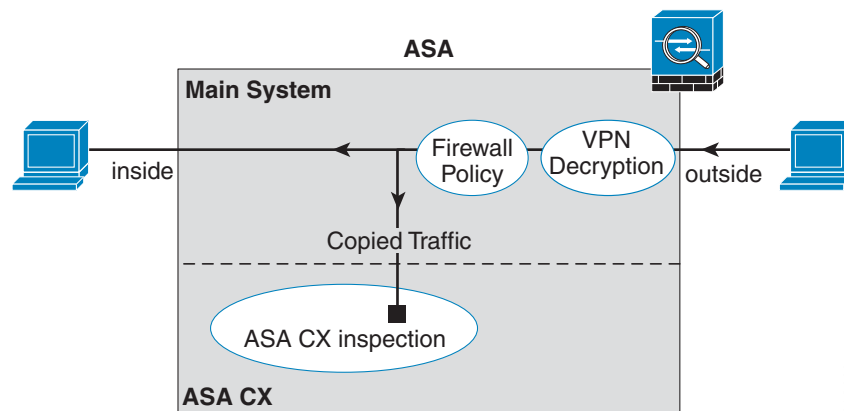
Traffic-Forwarding Interface in Monitor-Only Mode

You can alternatively configure ASA interfaces to be traffic-forwarding interfaces, where all traffic received is forwarded directly to the ASA CX module without any ASA processing. For testing and demonstration purposes, traffic-forwarding removes the extra complication of ASA processing. Traffic-forwarding is only supported in monitor-only mode, so the ASA CX module drops the traffic after inspecting it. Figure 27-3 shows the ASA GigabitEthernet 0/3 interface configured for traffic-forwarding. That interface is connected to a switch SPAN port so the ASA CX module can inspect all of the network traffic.

Figure 27-3   ASA CX Traffic-Forwarding
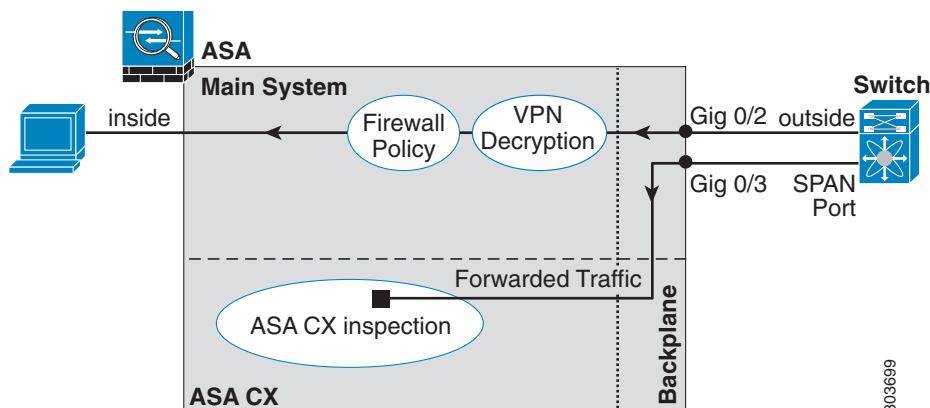
Information About ASA CX Management

- Initial Configuration, page 27-4
- Policy Configuration and Management, page 27-5

Initial Configuration

For initial configuration, you must use the CLI on the ASA CX module to run the **setup** command and configure other optional settings.

To access the CLI, you can use the following methods:

- ASA 5585-X:
  - ASA CX console port—The ASA CX console port is a separate external console port.
  - ASA CX Management 1/0 interface using SSH—You can connect to the default IP address (192.168.8.8), or you can use ASDM to change the management IP address and then connect using SSH. The ASA CX management interface is a separate external Gigabit Ethernet interface.

    ✎ **Note**   You cannot access the ASA CX hardware module CLI over the ASA backplane using the **session** command.

- ASA 5512-X through ASA 5555-X:
  - ASA session over the backplane—If you have CLI access to the ASA, then you can session to the module and access the module CLI.
  - ASA CX Management 0/0 interface using SSH—You can connect to the default IP address (192.168.1.2), or you can use ASDM to change the management IP address and then connect using SSH. These models run the ASA CX module as a software module. The ASA CX management interface shares the Management 0/0 interface with the ASA. Separate MAC addresses and IP addresses are supported for the ASA and ASA CX module. You must perform configuration of the ASA CX IP address within the ASA CX operating system (using the CLI

or ASDM). However, physical characteristics (such as enabling the interface) are configured on the ASA. You can remove the ASA interface configuration (specifically the interface name) to dedicate this interface as an ASA CX-only interface. This interface is management-only.

Policy Configuration and Management

After you perform initial configuration, configure the ASA CX policy using Cisco Prime Security Manager (PRSM). Then configure the ASA policy for sending traffic to the ASA CX module using ASDM or the ASA CLI.

Note   When using PRSM in multiple device mode, you can configure the ASA policy for sending traffic to the ASA CX module within PRSM, instead of using ASDM or the ASA CLI. Using PRSM lets you consolodate management to a single management system. However, PRSM has some limitations when configuring the ASA service policy; see the ASA CX user guide for more information.

Information About Authentication Proxy

When the ASA CX needs to authenticate an HTTP user (to take advantage of identity policies), you must configure the ASA to act as an authentication proxy: the ASA CX module redirects authentication requests to the ASA interface IP address/proxy port. By default, the port is 885 (user configurable). Configure this feature as part of the service policy to divert traffic from the ASA to the ASA CX module. If you do not enable the authentication proxy, only passive authentication is available.

Note   If you have a connection between hosts on two ASA interfaces, and the ASA CX service policy is only configured for one of the interfaces, then all traffic between these hosts is sent to the ASA CX module, including traffic orginiating on the non-ASA CX interface (the feature is bidirectional). However, the ASA only performs the authentication proxy on the interface to which the service policy is applied, because this feature is ingress-only.

Information About VPN and the ASA CX Module

The ASA includes VPN client and user authentication metadata from the Cisco AnyConnect client when forwarding traffic to the ASA CX module, which allows the ASA CX module to include this information as part of its policy lookup criteria. The VPN metadata is sent only at VPN tunnel establishment time along with a type-length-value (TLV) containing the session ID. The ASA CX module caches the VPN metadata for each session. Each tunneled connection sends the session ID so the ASA CX module can look up that session's metadata.

Compatibility with ASA Features

The ASA includes many advanced application inspection features, including HTTP inspection. However, the ASA CX module provides more advanced HTTP inspection than the ASA provides, as well as additional features for other applications, including monitoring and controlling application usage.

To take full advantage of the ASA CX module features, see the following guidelines for traffic that you send to the ASA CX module:

- Do not configure ASA inspection on HTTP traffic.
- Do not configure Cloud Web Security (ScanSafe) inspection. If you configure both the ASA CX action and Cloud Web Security inspection for the same traffic, the ASA only performs the ASA CX action.
- Other application inspections on the ASA are compatible with the ASA CX module, including the default inspections.
- Do not enable the Mobile User Security (MUS) server; it is not compatible with the ASA CX module.
- Do not enable ASA clustering; it is not compatible with the ASA CX module.
- If you enable failover, when the ASA fails over, any existing ASA CX flows are transferred to the new ASA, but the traffic is allowed through the ASA without being acted upon by the ASA CX module. Only new flows recieved by the new ASA are acted upon by the ASA CX module.
- (9.1(1) and earlier) Does not support NAT 64. In 9.1(2) and later, NAT 64 is supported.

Licensing Requirements for the ASA CX Module

Model	License Requirement
ASAv	Standard or Premium License.
All other models	Base License.

The ASA CX module and PRSM require additional licenses. See the ASA CX documentation for more information.

Prerequisites

To use PRSM to configure the ASA, you need to install a certificate on the ASA for secure communications. By default, the ASA generates a self-signed certificate. However, this certificate can cause browser prompts asking you to verify the certificate because the publisher is unknown. To avoid these browser prompts, you can instead install a certificate from a known certificate authority (CA). If you request a certificate from a CA, be sure the certificate type is both a server authentication certificate and a client authentication certificate. See the general operations configuration guide for more information.

Guidelines and Limitations

Context Mode Guidelines

(9.1(2) and earlier) Supported in single context mode only. Does not support multiple context mode.

(9.1(3) and later) Supported in multiple context mode. See the following guidelines:

- The ASA CX module itself (configured in PRSM) is a single context mode device; the context-specific traffic coming from the ASA is checked against the common ASA CX policy.

- For ASA CX module support, you cannot use the same IP addresses in multiple contexts; each context must include unique networks.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode. Traffic-forwarding interfaces are only supported in transparent mode.

Failover Guidelines

Does not support failover directly; when the ASA fails over, any existing ASA CX flows are transferred to the new ASA, but the traffic is allowed through the ASA without being inspected by the ASA CX.

ASA Clustering Guidelines

Does not support clustering.

IPv6 Guidelines

- Supports IPv6.

- (9.1(1) and earlier) Does not support NAT 64. In 9.1(2) and later, NAT 64 is supported.

Model Guidelines

- Supported only on the ASA 5585-X and 5512-X through ASA 5555-X. See the *Cisco ASA Compatibility Matrix* for more information:

  http://www.cisco.com/en/US/docs/security/asa/compatibility/asamatrx.html

- For the 5512-X through ASA 5555-X, you must install a Cisco solid state drive (SSD). For more information, see the ASA 5500-X hardware guide.

Monitor-Only Mode Guidelines

- You cannot configure both monitor-only mode and normal inline mode at the same time on the ASA. Only one type of security policy is allowed. In multiple context mode, you cannot configure monitor-only mode for some contexts, and regular inline mode for others.

- The following features are not supported in monitor-only mode:
  - Deny policies
  - Active authentication
  - Decryption policies

- The ASA CX does not perform packet buffering in monitor-only mode, and events will be generated on a best-effort basis. For example, some events, such as ones with long URLs spanning packet boundaries, may be impacted by the lack of buffering.

- Be sure to configure both the ASA policy and the ASA CX to have matching modes: both in monitor-only mode, or both in normal inline mode.

Additional guidelines for traffic-forwarding interfaces:

- The ASA must be in transparent mode.

- You can configure up to 4 interfaces as traffic-forwarding interfaces. Other ASA interfaces can be used as normal.

- Traffic-forwarding interfaces must be physical interfaces, not VLANs or BVIs. The physical interface also cannot have any VLANs associated with it.

- Traffic-forwarding interfaces cannot be used for ASA traffic; you cannot name them or configure them for ASA features, including failover or management-only.

- You cannot configure both a traffic-forwarding interface and a service policy for ASA CX traffic.

Additional Guidelines and Limitations

- See Compatibility with ASA Features, page 27-5.

- You cannot change the software type installed on the hardware module; if you purchase an ASA CX module, you cannot later install other software on it.

Default Settings

Table 27-1 lists the default settings for the ASA CX module.

Table 27-1   Default Network Parameters

Parameters	Default
Management IP address	ASA 5585-X: Management 1/0 192.168.8.8/24
	ASA 5512-X through ASA 5555-X: Management 0/0 192.168.1.2/24
Gateway	ASA 5585-X: 192.168.8.1/24
	ASA 5512-X through ASA 5555-X: 192.168.1.1/24
SSH or session Username	admin
Password	Admin123

Configuring the ASA CX Module

This section describes how to configure the ASA CX module.

Task Flow for the ASA CX Module

Configuring the ASA CX module is a process that includes configuration of the ASA CX security policy on the ASA CX module and then configuration of the ASA to send traffic to the ASA CX module. To configure the ASA CX module, perform the following steps:

Step 1 Cable the ASA CX management interfaces interface. See Connecting the ASA CX Management Interface, page 27-10.

Step 2 (ASA 5512-X through ASA 5555-X; May be required) Install the software module. See (ASA 5512-X through ASA 5555-X; May Be Required) Installing the Software Module, page 27-13.

Step 3 (ASA 5585-X) Configure the ASA CX module management IP address for initial SSH access. See (ASA 5585-X) Changing the ASA CX Management IP Address, page 27-14.

Step 4 On the ASA CX module, configure basic settings. You must use the CLI to configure these settings. See Configuring Basic ASA CX Settings at the ASA CX CLI, page 27-16.

Step 5 On the ASA CX module, configure the security policy using PRSM. See Configuring the Security Policy on the ASA CX Module Using PRSM, page 27-18.

Step 6 (Optional) On the ASA, configure the authentication proxy port. See (Optional) Configuring the Authentication Proxy Port, page 27-19.

Step 7 On the ASA, identify traffic to divert to the ASA CX module. See Redirecting Traffic to the ASA CX Module, page 27-19.

Note When using PRSM in multiple device mode, you can configure the ASA policy for sending traffic to the ASA CX module within PRSM, instead of using ASDM or the ASA CLI. However, PRSM has some limitations when configuring the ASA service policy; see the ASA CX user guide for more information.

Connecting the ASA CX Management Interface

In addition to providing management access to the ASA CX module, the ASA CX management interface needs access to an HTTP proxy server or a DNS server and the Internet for signature updates and more. This section describes recommended network configurations. Your network may differ.

ASA 5585-X (Hardware Module)

The ASA CX module includes a separate management interface from the ASA. For initial setup, you can connect with SSH to the ASA CX Management 1/0 interface using the default IP address (192.168.8.8/24). If you cannot use the default IP address, you can either use the console port or use ASDM to change the management IP address so you can use SSH.

If you have an inside router

If you have an inside router, you can route between the management network, which can include both the ASA Management 0/0 and ASA CX Management 1/0 interfaces, and the ASA inside network for Internet access. Be sure to also add a route on the ASA to reach the Management network through the inside router.
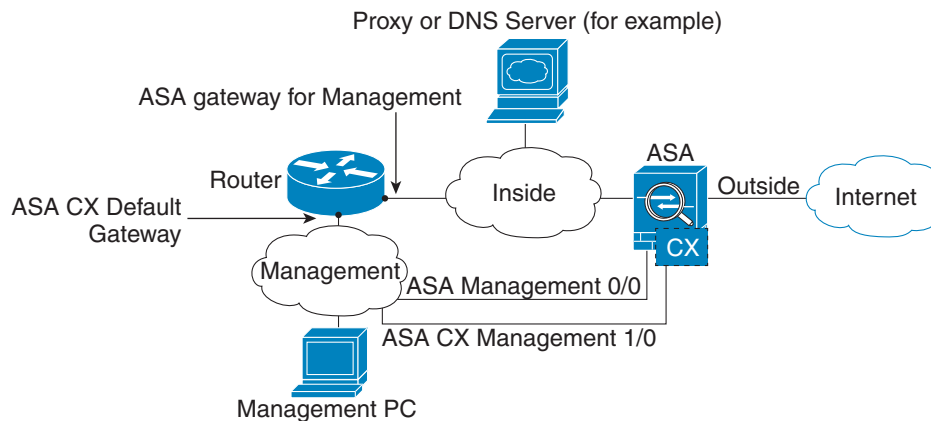
If you do not have an inside router

If you have only one inside network, then you cannot also have a separate management network, which would require an inside router to route between the networks. In this case, you can manage the ASA from the inside interface instead of the Management 0/0 interface. Because the ASA CX module is a separate device from the ASA, you can configure the ASA CX Management 1/0 address to be on the same network as the inside interface.
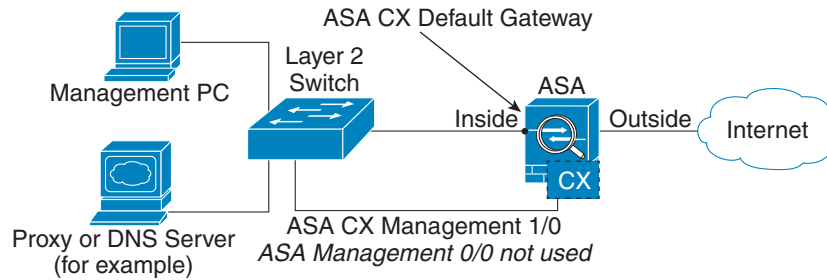
ASA 5512-X through ASA 5555-X (Software Module)

These models run the ASA CX module as a software module, and the ASA CX management interface shares the Management 0/0 interface with the ASA.

If you have an inside router

If you have an inside router, you can route between the Management 0/0 network, which includes both the ASA and ASA CX management IP addresses, and the inside network for Internet access. Be sure to also add a route on the ASA to reach the Management network through the inside router.
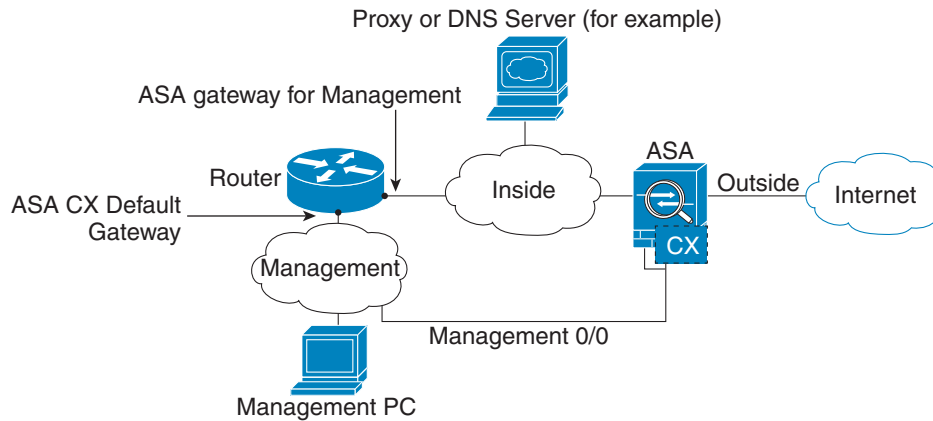
If you do not have an inside router

If you have only one inside network, then you cannot also have a separate management network. In this case, you can manage the ASA from the inside interface instead of the Management 0/0 interface. If you remove the ASA-configured name from the Management 0/0 interface, you can still configure the ASA CX IP address for that interface. Because the ASA CX module is essentially a separate device from the ASA, you *can* configure the ASA CX management address to be on the same network as the inside interface.
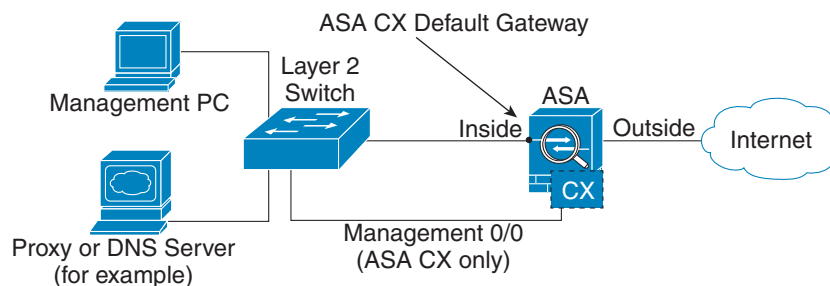
Note   You must remove the ASA-configured name for Management 0/0; if it is configured on the ASA, then the ASA CX address must be on the same network as the ASA, and that excludes any networks already configured on other ASA interfaces. If the name is not configured, then the ASA CX address can be on any network, for example, the ASA inside network.

What to Do Next

- Configure the ASA CX management IP address. See .

(ASA 5512-X through ASA 5555-X; May Be Required) Installing the Software Module

If you purchase the ASA with the ASA CX module, the module software and required solid state drive(s) (SSDs) come pre-installed and ready to go. If you want to add the ASA CX to an existing ASA, or need to replace the SSD, you need to install the ASA CX boot software and partition the SSD according to this procedure. To physically install the SSD, see the ASA hardware guide.

> **Note**   For the ASA 5585-X hardware module, you must install or upgrade your image from within the ASA CX module. See the ASA CX module documentation for more information.

Prerequisites

- The free space on flash (disk0) should be at least 3GB plus the size of the boot software.
- In multiple context mode, perform this procedure in the system execution space.

Detailed Steps

Step 1   Download the ASA CX boot software from Cisco.com to your computer. If you have a Cisco.com login, you can obtain the boot software from the following website:

http://www.cisco.com/cisco/software/release.html?mdfid=284325223&softwareid=284399946

The boot software lets you set basic ASA CX network configuration, partition the SSD, and download the larger system software from a server of your choice to the SSD.

Step 2   Download the ASA CX system software from Cisco.com to an HTTP, HTTPS, or FTP server accessible from the ASA CX management interface. If you have a Cisco.com login, you can obtain the boot software from the following website:

http://www.cisco.com/cisco/software/release.html?mdfid=284325223&softwareid=284399946

Step 3   In ASDM, choose **Tools > File Management**, and then choose **File Transfer > Between Local PC and Flash**. Transfer the boot software to disk0 on the ASA. Do not transfer the system software; it is downloaded later to the SSD.

Step 4   Connect to the ASA CLI, and enter privileged EXEC mode. See "Getting Started" chapter in the general operations configuration guide to access the ASA CLI.

Step 5   If you are replacing the IPS module with the ASA CX module, shut down and uninstall the IPS module, and then reload the ASA:

```
hostname# sw-module module ips shutdown
hostname# sw-module module ips uninstall
hostname# reload
```

After the ASA reloads, reconnect to the ASA CLI.

Step 6   Set the ASA CX module boot image location in ASA disk0 by entering the following command:

```
hostname# sw-module module cxsc recover configure image disk0:file_path
```

Example:

```
hostname# sw-module module cxsc recover configure image disk0:asacx-boot-9.1.1.img
```

Step 7   Load the ASA CX boot image by entering the following command:

```
hostname# sw-module module cxsc recover boot
```

Step 8   Wait approximately 5 minutes for the ASA CX module to boot up, and then open a console session to the now-running ASA CX boot image. The default username is **admin** and the default password is **Admin123**.

```
hostname# session cxsc console
Establishing console session with slot 1
Opening console session with module cxsc.
Connected to module cxsc. Escape character sequence is 'CTRL-SHIFT-6 then x'.
cxsc login: admin
Password: Admin123
```

Step 9   Partition the SSD:

```
asacx-boot> partition
....
Partition Successfully Completed
```

Step 10 Perform the basic network setup using the **setup** command according to the Configuring Basic ASA CX Settings at the ASA CX CLI, page 27-16 (do not exit the ASA CX CLI), and then return to this procedure to install the software image.

Step 11 Install the system software from the server:

```
asacx-boot> system install url
```

Example:

The following command installs the asacx-sys-9.1.1.pkg system software.

```
asacx-boot> system install https://upgrades.example.com/packages/asacx-sys-9.1.1.pkg

Username: buffy
Password: angelforever
Verifying
Downloading
Extracting
Package Detail
        Description:
        Requires reboot:
Cisco ASA CX System Upgrade
Yes
Do you want to continue with upgrade? [n]: Y
Warning: Please do not interrupt the process or turn off the system. Doing so might leave
system in unusable state.
Upgrading
Stopping all the services ...
Starting upgrade process ...
Reboot is required to complete the upgrade. Press Enter to reboot the system.
```

Step 12 Press **Enter** to reboot the ASA CX module. Rebooting the module closes the console session. Allow 10 or more minutes for application component installation and for the ASA CX services to start.

(ASA 5585-X) Changing the ASA CX Management IP Address

If you cannot use the default management IP address (192.168.8.8), then you can set the management IP address from the ASA. After you set the management IP address, you can access the ASA CX module using SSH to perform initial setup.

> Note   For a software module, you can access the ASA CX CLI to perform setup by sessioning from the ASA
> CLI; you can then set the ASA CX management IP address as part of setup. See Configuring Basic ASA
> CX Settings at the ASA CX CLI, page 27-16.

Guidelines

In multiple context mode, perform this procedure in the system execution space.

Detailed Steps

Multiple Context Mode

Step 1   In the System, choose Tools > Command Line Interface.

Step 2   Enter the following command:

Command	Purpose

`session 1 do setup host ip`
`ip_address/mask,gateway_ip`
 	
Example:
`hostname# session 1 do setup host ip`
`10.1.1.2/24,10.1.1.1`
	Sets the ASA CX management IP address, mask, and gateway.

Step 3   Click Send.

Single Context Mode

Step 1   In ASDM, choose **Wizards > Startup Wizard**.

Step 2   Click **Next** to advance through the initial screens until you reach the ASA CX Basic Configuration
screen.

Step 3   Enter the new management IP address, subnet mask, and default gateway.

Step 4   (Optional) Change the Auth Proxy Port. You can set this later if desired. See (Optional) Configuring the Authentication Proxy Port, page 27-19 for more information.

Step 5   Click **Finish** to skip the remaining screens, or click **Next** to advance through the remaining screens and complete the wizard.

Configuring Basic ASA CX Settings at the ASA CX CLI

You must configure basic network settings and other parameters on the ASA CX module before you can configure your security policy.

Detailed Steps

Step 1   Do one of the following:

- (All models) Use SSH to connect to the ASA CX management IP address.

- (ASA 5512-X through ASA 5555-X) Open a console session to the module from the ASA CLI (see the "Getting Started" chapter in the general operations configuration guide to access the ASA CLI). In multiple context mode, session from the system execution space.

```
hostname# session cxsc console
```

Step 2   Log in with the username **admin** and the password **Admin123**. You will change the password as part of this procedure.

Step 3   Enter the following command:

```
asacx> setup
```

Example:

```
asacx> setup
Welcome to Cisco Prime Security Manager Setup
[hit Ctrl-C to abort]
Default values are inside [ ]
```

You are prompted through the setup wizard. The following example shows a typical path through the wizard; if you enter **Y** instead of **N** at a prompt, you will be able to configure some additional settings. This example shows how to configure both IPv4 and IPv6 static addresses. You can configure IPv6 stateless auto configuration by answering **N** when asked if you want to configure a static IPv6 address.

```
Enter a hostname [asacx]: asa-cx-host
Do you want to configure IPv4 address on management interface?(y/n) [Y]: Y
Do you want to enable DHCP for IPv4 address assignment on management interface?(y/n)[N]: N
Enter an IPv4 address [192.168.8.8]: 10.89.31.65
Enter the netmask [255.255.255.0]: 255.255.255.0
Enter the gateway [192.168.8.1]: 10.89.31.1
Do you want to configure static IPv6 address on management interface?(y/n) [N]: Y
Enter an IPv6 address: 2001:DB8:0:CD30::1234/64
Enter the gateway: 2001:DB8:0:CD30::1
Enter the primary DNS server IP address [ ]: 10.89.47.11
Do you want to configure Secondary DNS Server? (y/n) [N]: N
Do you want to configure Local Domain Name? (y/n) [N] Y
Enter the local domain name: example.com
Do you want to configure Search domains? (y/n) [N] Y
Enter the comma separated list for search domains: example.com
Do you want to enable the NTP service?(y/n) [N]: Y
Enter the NTP servers separated by commas: 1.ntp.example.com, 2.ntp.example.com
```

Step 4   After you complete the final prompt, you are presented with a summary of the settings. Look over the summary to verify that the values are correct, and enter **Y** to apply your changed configuration. Enter **N** to cancel your changes.

Example:

```
Apply the changes?(y,n) [Y]: Y
Configuration saved successfully!
Applying...
Done.
Generating self-signed certificate, the web server will be restarted after that
...
Done.
Press ENTER to continue...
asacx>
```

🖉

Note   If you change the host name, the prompt does not show the new name until you log out and log back in.

Step 5   If you do not use NTP, configure the time settings. The default time zone is the UTC time zone. Use the **show time** command to see the current settings. You can use the following commands to change time settings:

```
asacx> config timezone
asacx> config time
```

Step 6 Change the admin password by entering the following command:

```
asacx> config passwd
```

Example:

```
asacx> config passwd
The password must be at least 8 characters long and must contain
at least one uppercase letter (A-Z), at least one lowercase letter
(a-z) and at least one digit (0-9).
Enter password: Farscape1
Confirm password: Farscape1
SUCCESS: Password changed for user admin
```

Step 7 Enter the **exit** command to log out.

Configuring the Security Policy on the ASA CX Module Using PRSM

This section describes how to launch PRSM to configure the ASA CX module application. For details on using PRSM to configure your ASA CX security policy, see the ASA CX user guide.

Detailed Steps

You can launch PRSM from your web browser, or you can launch it from ASDM.

- Launch PRSM from a web browser by enter the following URL:

  https://*ASA_CX_management_IP*

  Where the ASA CX management IP address is the one you set in the Configuring Basic ASA CX Settings at the ASA CX CLI, page 27-16.

- Launch PRSM from ASDM by choosing **Home > ASA CX Status**, and clicking the **Connect to the ASA CX application** link.

What to Do Next

- (Optional) Configure the authentication proxy port. See (Optional) Configuring the Authentication Proxy Port, page 27-19.
- Redirect traffic to the ASA CX module. See Redirecting Traffic to the ASA CX Module, page 27-19.

(Optional) Configuring the Authentication Proxy Port

The default authentication port is 885. To change the authentication proxy port, perform the following steps. For more information about the authentication proxy, see Information About Authentication Proxy, page 27-5.

✎
Note   (Single mode) You can also set the port as part of the ASDM startup wizard. See (ASA 5585-X) Changing the ASA CX Management IP Address, page 27-14.

Guidelines

In multiple context mode, perform this procedure within each security context.

Detailed Steps

Step 1   In ASDM, choose **Configuration > Firewall > Advanced > ASA CX Auth Proxy**.

○ ○ **Configuration > Firewall > Advanced > ASA CX Auth Proxy**

Configure ASA CX auth proxy port.

Auth proxy port value (must be greater than 1024): [5000]

Step 2   Enter a port greater than 1024. The default is 885.

Step 3   Click **Apply**.

Redirecting Traffic to the ASA CX Module

You can redirect traffic to the ASA CX module by creating a service policy that identifies specific traffic. For demonstration purposes only, you can also enable monitor-only mode for the service policy, which forwards a copy of traffic to the ASA CX module, while the original traffic remains unaffected.

Another option for demonstration purposes is to configure a traffic-forwarding interface instead of a service policy in monitor-only mode. The traffic-forwarding interface sends all traffic directly to the ASA CX module, bypassing the ASA.

- Creating the ASA CX Service Policy, page 27-20
- Configuring Traffic-Forwarding Interfaces (Monitor-Only Mode), page 27-22

Creating the ASA CX Service Policy

This section identifies traffic to redirect from the ASA to the ASA CX module. Configure this policy on the ASA. If you want to use a traffic-forwarding interface for demonstration purposes, skip this procedure and see Configuring Traffic-Forwarding Interfaces (Monitor-Only Mode), page 27-22 instead.

✎ **Note**   When using PRSM in multiple device mode, you can configure the ASA policy for sending traffic to the ASA CX module within PRSM, instead of using ASDM or the ASA CLI. However, PRSM has some limitations when configuring the ASA service policy; see the ASA CX user guide for more information.

Prerequisites

- If you enable the authentication proxy on the ASA using this procedure, be sure to also configure a directory realm for authentication on the ASA CX module. See the ASA CX user guide for more information.

- If you have an active service policy redirecting traffic to an IPS module (that you replaced with the ASA CX), you must remove that policy before you configure the ASA CX service policy.

- Be sure to configure both the ASA policy and the ASA CX to have matching modes: both in monitor-only mode, or both in normal inline mode.

- In multiple context mode, perform this procedure within each security context.

Detailed Steps

Step 1   Choose **Configuration > Firewall > Service Policy Rules**.

Step 2   Choose **Add > Add Service Policy Rule**. The Add Service Policy Rule Wizard - Service Policy dialog box appears.

Step 3   Complete the Service Policy dialog box as desired. See the ASDM online help for more information about these screens.

Step 4   Click **Next**. The Add Service Policy Rule Wizard - Traffic Classification Criteria dialog box appears.

Step 5   Complete the Traffic Classification Criteria dialog box as desired. See the ASDM online help for more information about these screens.

Step 6   Click **Next** to show the Add Service Policy Rule Wizard - Rule Actions dialog box.

Step 7   Click the **ASA CX Inspection** tab.

Step 8   Check the **Enable ASA CX for this traffic flow** check box.

Step 9   In the If ASA CX Card Fails area, click one of the following:

- **Permit traffic**—Sets the ASA to allow all traffic through, uninspected, if the ASA CX module is unavailable.

- **Close traffic**—Sets the ASA to block all traffic if the ASA CX module is unavailable.

Step 10   (Optional) To enable the authentication proxy, which is required for active authentication, check the **Enable Auth Proxy** check box. This option is not available in monitor-only mode.

Step 11   (Optional) For demonstration purposes only, check the **Monitor-only** check box to send a read-only copy of traffic to the ASA CX module. See Monitor-Only Mode, page 27-3 for more information.

✎

Note   You must configure all classes and policies to be either in monitor-only mode, or in normal inline mode; you cannot mix both modes on the same ASA.

Step 12   Click **OK** and then **Apply**.

Step 13   Repeat this procedure to configure additional traffic flows as desired.

Configuring Traffic-Forwarding Interfaces (Monitor-Only Mode)

This section configures traffic-forwarding interfaces, where all traffic is forwarded directly to the ASA CX module. This method is for demonstration purposes only. For a normal ASA CX service policy, see Creating the ASA CX Service Policy, page 27-20.

For more information see Monitor-Only Mode, page 27-3. See also the Guidelines and Limitations, page 27-6 for guidelines and limitations specific to traffic-forwarding interfaces.

You can only configure this feature at the CLI; you can use the Command Line Interface tool.

Prerequisites

- Be sure to configure both the ASA policy and the ASA CX to have matching modes: both in monitor-only.
- In multiple context mode, perform this procedure within each security context.

Detailed Steps

Step 1   Choose **Tools > Command Line Interface**.

Step 2   Click the **Multiple Line** radio button.

Step 3   Enter the following commands:

	Command	Purpose
Step 1	**interface** *physical_interface* **Example:** hostname(config)# interface gigabitethernet 0/5	Enters interface configuration mode for the physical interface you want to use for traffic-forwarding.
Step 2	**no nameif** **Example:** hostname(config-ifc)# no nameif	Removes any name configured for the interface. If this interface was used in any ASA configuration, that configuration is removed. You cannot configure traffic-forwarding on a named interface.
Step 3	**traffic-forward cxsc monitor-only** **Example:** hostname(config-ifc)# traffic-forward cxsc monitor-only	Enables traffic-forwarding. You see a warning similar to the following: WARNING: This configuration is purely for demo of CX functionality and shouldn't be used on a production ASA and any issues found when mixing demo feature with production ASA is not supported.
Step 4	**no shutdown** **Example:** hostname(config-ifc)# no shutdown	Enables the interface.

Step 4   Repeat for any additional interfaces.

Step 5 Click **Send**.

Examples

The following example makes GigabitEthernet 0/5 a traffic-forwarding interface:
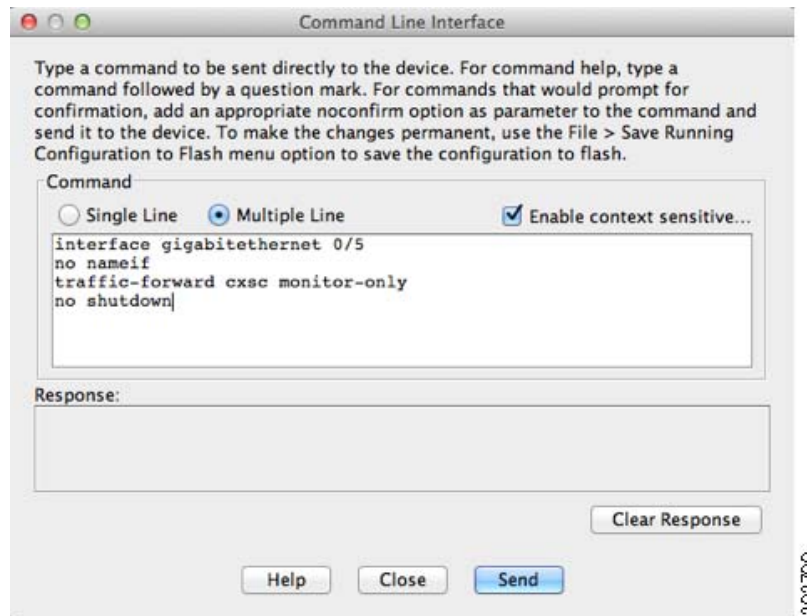
Managing the ASA CX Module

This section includes procedures that help you manage the module.

Resetting the Password

You can reset the module password to the default. For the user **admin**, the default password is **Admin123**. After resetting the password, you should change it to a unique value using the module application.

Resetting the module password causes the module to reboot. Services are not available while the module is rebooting.

If you cannot connect to ASDM with the new password, restart ASDM and try to log in again. If you defined a new password and still have an existing password in ASDM that is different from the new password, clear the password cache by choosing **File > Clear ASDM Password Cache**, then restart ASDM and try to log in again.

To reset the module password to the default of Admin123, perform the following steps.

Guidelines

In multiple context mode, perform this procedure in the system execution space.

Detailed Steps

Step 1   From the ASDM menu bar, choose **Tools > ASA CX Password Reset**.

The Password Reset confirmation dialog box appears.

Step 2   Click **OK** to reset the password to the default **Admin123**.

A dialog box displays the success or failure of the password reset.

Step 3   Click **Close** to close the dialog box.

Reloading or Resetting the Module

To reload or reset the module, enter one of the following commands at the ASA CLI.

Guidelines

In multiple context mode, perform this procedure in the system execution space.

Detailed Steps

Command	Purpose
For a hardware module (ASA 5585-X): **hw-module module 1 reload** For a software module (ASA 5512-X through ASA 5555-X): **sw-module module cxsc reload** **Example:** hostname# hw-module module 1 reload	Reloads the module software.
For a hardware module: **hw-module module 1 reset** For a software module: **sw-module module cxsc reset** **Example:** hostname# hw-module module 1 reset	Performs a reset, and then reloads the module.

Shutting Down the Module

Shutting down the module software prepares the module to be safely powered off without losing configuration data. **Note**: If you reload the ASA, the module is not automatically shut down, so we recommend shutting down the module before reloading the ASA. To gracefully shut down the module, perform the following steps at the ASA CLI.

Guidelines

In multiple context mode, perform this procedure in the system execution space.

Detailed Steps

Command	Purpose
For a hardware module (ASA 5585-X): **hw-module module 1 shutdown** For a software module (ASA 5512-X through ASA 5555-X): **sw-module module cxsc shutdown** **Example:** hostname# hw-module module 1 shutdown	Shuts down the module.

(ASA 5512-X through ASA 5555-X) Uninstalling a Software Module Image

To uninstall a software module image and associated configuration, perform the following steps.

Guidelines

In multiple context mode, perform this procedure in the system execution space.

Detailed Steps

	Command	Purpose

Step 1	`sw-module module cxsc uninstall`
	
	Example:
	`hostname# sw-module module cxsc uninstall`
	`Module cxsc will be uninstalled. This will completely remove the`
	`disk image associated with the sw-module including any configuration`
	`that existed within it.`
	
	`Uninstall module <id>? [confirm]`
		Permanently uninstalls the software module image and associated configuration.

Step 2	`reload`
	
	Example:
	`hostname# reload`
		Reloads the ASA. You must reload the ASA before you can install a new module type.

(ASA 5512-X through ASA 5555-X) Sessioning to the Module From the ASA

To access the ASA CX software module CLI from the ASA, you can session from the ASA. You can either session to the module (using Telnet) or create a virtual console session. A console session might be useful if the control plane is down and you cannot establish a Telnet session.

You may need to access the CLI if you are using multiple context mode and you need to set basic network settings using the CLI, or for troubleshooting.

Guidelines

In multiple context mode, perform this procedure in the system execution space.

Detailed Steps

Command	Purpose

Telnet session.
 	
session cxsc
 	
Example:
`hostname# session cxsc`
 	
`Opening command session with slot 1.`
`Connected to module cxsc. Escape character`
`sequence is 'CTRL-^X'.`
 	
`cxsc login: admin`
`Password: Admin123`
	Accesses the module using Telnet. You are prompted for the username and password. The default username is **admin**, and the default password is **Admin123**.

Console session.
 	
session cxsc console
 	
Example:
`hostname# session cxsc console`
 	
`Establishing console session with slot 1`
`Opening console session with module cxsc.`
`Connected to module cxsc. Escape character`
`sequence is 'CTRL-SHIFT-6 then x'.`
 	
`cxsc login: admin`
`Password: Admin123`
	Accesses the module console. You are prompted for the username and password. The default username is **admin**, and the default password is **Admin123**.
	
	Note Do not use this command in conjunction with a terminal server where **Ctrl-Shift-6, x** is the escape sequence to return to the terminal server prompt. **Ctrl-Shift-6, x** is also the sequence to escape the ASA CX console and return to the ASA prompt. Therefore, if you try to exit the ASA CX console in this situation, you instead exit all the way to the terminal server prompt. If you reconnect the terminal server to the ASA, the ASA CX console session is still active; you can never exit to the ASA prompt. You must use a direct serial connection to return the console to the ASA prompt.
	
	Use the **session cxsc** command instead.

Monitoring the ASA CX Module

Use Tools > Command Line Interface to use monitoring commands.

Note   For ASA CX-related syslog messages, see the syslog messages guide. ASA CX syslog messages start with message number 429001.

Showing Module Status

See the ASA CX Status Tab in the general operations configuration guide.

Showing Module Statistics

To show module statistics, enter the following command:

Command	Purpose
`show service-policy cxsc`	Displays the ASA CX statistics and status per service policy.

Examples

The following is sample output from the **show service-policy** command showing the ASA CX policy and the current statistics as well as the module status when the authentication proxy is disabled:

```
hostname# show service-policy cxsc
Global policy:
  Service-policy: global_policy
    Class-map: bypass
      CXSC: card status Up, mode fail-open, auth-proxy disabled
        packet input 2626422041, packet output 2626877967, drop 0, reset-drop 0, proxied 0
```

The following is sample output from the **show service-policy** command showing the ASA CX policy and the current statistics as well as the module status when the authentication proxy is enabled; in this case, the proxied counters also increment:

```
hostname# show service-policy cxsc
Global policy:
  Service-policy: pmap
    Class-map: class-default
      Default Queueing      Set connection policy: random-sequence-number disable
        drop 0
      CXSC: card status Up, mode fail-open, auth-proxy enabled
        packet input 7724, packet output 7701, drop 0, reset-drop 0, proxied 10
```

Monitoring Module Connections

To show connections through the ASA CX module, enter one of the following commands:

Command	Purpose
`show asp table classify domain cxsc`	Shows the NP rules created to send traffic to the ASA CX module.
`show asp table classify domain cxsc-auth-proxy`	Shows the NP rules created for the authentication proxy for the ASA CX module.

Command	Purpose
`show asp drop`	Shows dropped packets. The following drop types are used: Frame Drops: • cxsc-bad-tlv-received—This occurs when ASA receives a packet from CXSC without a Policy ID TLV. This TLV must be present in non-control packets if it does not have the Standy Active bit set in the actions field. • cxsc-request—The frame was requested to be dropped by CXSC due a policy on CXSC whereby CXSC would set the actions to Deny Source, Deny Destination, or Deny Pkt. • cxsc-fail-close—The packet is dropped because the card is not up and the policy configured was 'fail-close' (rather than 'fail-open' which allows packets through even if the card was down). • cxsc-fail—The CXSC configuration was removed for an existing flow and we are not able to process it through CXSC it will be dropped. This should be very unlikely. • cxsc-malformed-packet—The packet from CXSC contains an invalid header. For instance, the header length may not be correct. Flow Drops: • cxsc-request—The CXSC requested to terminate the flow. The actions bit 0 is set. • reset-by-cxsc—The CXSC requested to terminate and reset the flow. The actions bit 1 is set. • cxsc-fail-close—The flow was terminated because the card is down and the configured policy was 'fail-close'.
`show asp event dp-cp cxsc-msg`	This output shows how many ASA CX module messages are on the dp-cp queue. Currently, only VPN queries from the ASA CX module are sent to dp-cp.
`show conn`	This command already shows if a connection is being forwarded to a module by displaying the 'X - inspected by service module' flag. Connections being forwarded to the ASA CX module will also display the 'X' flag.

Examples

The following is sample output from the **show asp table classify domain cxsc** command:

```
hostname# show asp table classify domain cxsc
Input Table
in  id=0x7ffedb4acf40, priority=50, domain=cxsc, deny=false
    hits=15485658, user_data=0x7ffedb4ac840, cs_id=0x0, use_real_addr, flags=0x0,
protocol=0
    src ip/id=0.0.0.0, mask=0.0.0.0, port=0
    dst ip/id=0.0.0.0, mask=0.0.0.0, port=0, dscp=0x0
    input_ifc=outside, output_ifc=any
in  id=0x7ffedb4ad4a0, priority=50, domain=cxsc, deny=false
    hits=992053, user_data=0x7ffedb4ac840, cs_id=0x0, use_real_addr, flags=0x0, protocol=0
    src ip/id=0.0.0.0, mask=0.0.0.0, port=0
    dst ip/id=0.0.0.0, mask=0.0.0.0, port=0, dscp=0x0
    input_ifc=inside, output_ifc=any
```

```
    in  id=0x7ffedb4ada00, priority=50, domain=cxsc, deny=false
        hits=0, user_data=0x7ffedb4ac840, cs_id=0x0, use_real_addr, flags=0x0, protocol=0
        src ip/id=0.0.0.0, mask=0.0.0.0, port=0
        dst ip/id=0.0.0.0, mask=0.0.0.0, port=0, dscp=0x0
        input_ifc=m, output_ifc=any

Output Table:

L2 - Output Table:

L2 - Input Table:

Last clearing of hits counters: Never
```

The following is sample output from the **show asp table classify domain cxsc-auth-proxy** command.
For the first rule in the output, the destination "port=2000" is the auth-proxy port configured by the **cxsc
auth-proxy port 2000** command, and the destination "ip/id=192.168.0.100" is the ASA interface IP
address.

```
hostname# show asp table classify domain cxsc-auth-proxy
Input Table
in  id=0x7ffed86cc470, priority=121, domain=cxsc-auth-proxy, deny=false
    hits=0, user_data=0x7ffed86ca220, cs_id=0x0, flags=0x0, protocol=6
    src ip/id=0.0.0.0, mask=0.0.0.0, port=0
    dst ip/id=192.168.0.100, mask=255.255.255.255, port=2000, dscp=0x0
    input_ifc=inside, output_ifc=identity
in  id=0x7ffed86cce20, priority=121, domain=cxsc-auth-proxy, deny=false
    hits=0, user_data=0x7ffed86ca220, cs_id=0x0, flags=0x0, protocol=6
    src ip/id=0.0.0.0, mask=0.0.0.0, port=0
    dst ip/id=2.2.2.2, mask=255.255.255.255, port=2000, dscp=0x0
    input_ifc=new2, output_ifc=identity
in  id=0x7ffed86cd7d0, priority=121, domain=cxsc-auth-proxy, deny=false
    hits=0, user_data=0x7ffed86ca220, cs_id=0x0, flags=0x0, protocol=6
    src ip/id=0.0.0.0, mask=0.0.0.0, port=0
    dst ip/id=172.23.58.52, mask=255.255.255.255, port=2000, dscp=0x0
    input_ifc=mgmt, output_ifc=identity
in  id=0x7ffed86caa80, priority=121, domain=cxsc-auth-proxy, deny=false
    hits=0, user_data=0x7ffed86ca220, cs_id=0x0, flags=0x0, protocol=6
    src ip/id=0.0.0.0, mask=0.0.0.0, port=0
    dst ip/id=192.168.5.172, mask=255.255.255.255, port=2000, dscp=0x0
    input_ifc=outside, output_ifc=identity
in  id=0x7ffed86cb3c0, priority=121, domain=cxsc-auth-proxy, deny=false
    hits=0, user_data=0x7ffed86ca220, cs_id=0x0, flags=0x0, protocol=6
    src ip/id=::/0, port=0
        dst ip/id=fe80::5675:d0ff:fe5b:1102/128, port=2000
    input_ifc=outside, output_ifc=identity
in  id=0x7ffed742be10, priority=121, domain=cxsc-auth-proxy, deny=false
    hits=0, user_data=0x7ffed86ca220, cs_id=0x0, flags=0x0, protocol=6
    src ip/id=::/0, port=0
    dst ip/id=1:1:1:1::10/128, port=2000
    input_ifc=outside, output_ifc=identity

Output Table:

L2 - Output Table:

L2 - Input Table:

Last clearing of hits counters: Never
```

The following is sample output from the **show asp drop** command. This output is just an example and
lists all the possible reasons for a dropped frame or flow from the ASA CX module:

```
hostname# show asp drop
Frame drop:
  CXSC Module received packet with bad TLV's (cxsc-bad-tlv-received)        2
  CXSC Module requested drop (cxsc-request)                                1
  CXSC card is down (cxsc-fail-close)                                      1
  CXSC config removed for flow (cxsc-fail)                                 3
  CXSC Module received malformed packet (cxsc-malformed-packet)            1

Last clearing: 18:12:58 UTC May 11 2012 by enable_15

Flow drop:
  Flow terminated by CXSC (cxsc-request)                                   2
  Flow reset by CXSC (reset-by-cxsc)                                       1
  CXSC fail-close (cxsc-fail-close)                                        1

Last clearing: 18:12:58 UTC May 11 2012 by enable_15
```

The following is sample output from the **show asp event dp-cp cxsc-msg** command:

```
hostname# show asp event dp-cp cxsc-msg
DP-CP EVENT QUEUE                 QUEUE-LEN  HIGH-WATER
Punt Event Queue                         0          5
Identity-Traffic Event Queue             0          0
General Event Queue                      0          4
Syslog Event Queue                       4         90
Non-Blocking Event Queue                 0          2
Midpath High Event Queue                 0         53
Midpath Norm Event Queue              8074       8288
SRTP Event Queue                         0          0
HA Event Queue                           0          0
Threat-Detection Event Queue             0          3
ARP Event Queue                          0       2048
IDFW Event Queue                         0          0
CXSC Event Queue                         0          1
EVENT-TYPE        ALLOC ALLOC-FAIL ENQUEUED ENQ-FAIL  RETIRED 15SEC-RATE
cxsc-msg              1         0        1        0        1          0
```

The following is sample output from the **show conn detail** command:

```
hostname# show conn detail
0 in use, 105 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module

TCP outside 208.80.152.2:80 inside 192.168.1.20:59928, idle 0:00:10, bytes 79174, flags
XUIO
```

Capturing Module Traffic

To configure and view packet captures for the ASA CX module, enter one of the following commands:

Command	Purpose
`capture` *name* `interface asa_dataplane`	Captures packets between ASA CX module and the ASA on the backplane.
`copy capture`	Copies the capture file to a server.
`show capture`	Shows the capture at the ASA console.

Note Captured packets contain an additional AFBP header that your PCAP viewer might not understand; be sure to use the appropriate plugin to view these packets.

Troubleshooting the ASA CX Module

- Problems with the Authentication Proxy, page 27-32

Problems with the Authentication Proxy

If you are having a problem using the authentication proxy feature, follow these steps to troubleshoot your configuration and connections:

1. Check your configurations.

   - On the ASA, check the output of the **show asp table classify domain cxsc-auth-proxy** command and make sure there are rules installed and that they are correct.

   - In PRSM, ensure the directory realm is created with the correct credentials and test the connection to make sure you can reach the authentication server; also ensure that a policy object or objects are configured for authentication.

2. Check the output of the **show service-policy cxsc** command to see if any packets were proxied.

3. Perform a packet capture on the backplane, and check to see if traffic is being redirected on the correct configured port. See Capturing Module Traffic, page 27-32. You can check the configured port using the **show running-config cxsc** command or the **show asp table classify domain cxsc-auth-proxy** command.

Note If you have a connection between hosts on two ASA interfaces, and the ASA CX service policy is only configured for one of the interfaces, then all traffic between these hosts is sent to the ASA CX module, including traffic orginiating on the non-ASA CX interface (the feature is bidirectional). However, the ASA only performs the authentication proxy on the interface to which the service policy is applied, because this feature is ingress-only.

Example 27-1   Make sure port 2000 is used consistently:

1. Check the authentication proxy port:

```
hostname# show running-config cxsc
cxsc auth-proxy port 2000
```

2. Check the authentication proxy rules:

```
hostname# show asp table classify domain cxsc-auth-proxy

Input Table
in  id=0x7ffed86cc470, priority=121, domain=cxsc-auth-proxy, deny=false
    hits=0, user_data=0x7ffed86ca220, cs_id=0x0, flags=0x0, protocol=6
    src ip/id=0.0.0.0, mask=0.0.0.0, port=0
    dst ip/id=192.168.0.100, mask=255.255.255.255, port=2000, dscp=0x0
    input_ifc=inside, output_ifc=identity
```

3. In the packet captures, the redirect request should be going to destination port 2000.

Feature History for the ASA CX Module

Table 27-2 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 27-2   Feature History for the ASA CX Module

Feature Name	Platform Releases	Feature Information
ASA 5585-X with SSP-10 and -20 support for the ASA CX SSP-10 and -20	ASA 8.4(4.1) ASA CX 9.0(1)	The ASA CX module lets you enforce security based on the complete context of a situation. This context includes the identity of the user (who), the application or website that the user is trying to access (what), the origin of the access attempt (where), the time of the attempted access (when), and the properties of the device used for the access (how). With the ASA CX module, you can extract the full context of a flow and enforce granular policies such as permitting access to Facebook but denying access to games on Facebook or permitting finance employees access to a sensitive enterprise database but denying the same access to other employees. We introduced the following screens: Home > ASA CX Status Wizards > Startup Wizard > ASA CX Basic Configuration Configuration > Firewall > Service Policy Rules > Add Service Policy Rule > Rule Actions > ASA CX Inspection
ASA 5512-X through ASA 5555-X support for the ASA CX SSP	ASA 9.1(1) ASA CX 9.1(1)	We introduced support for the ASA CX SSP software module for the ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, and ASA 5555-X. We did not modify any screens.

Table 27-2   Feature History for the ASA CX Module (continued)

Feature Name	Platform Releases	Feature Information
Monitor-only mode for demonstration purposes	ASA 9.1(2) ASA CX 9.1(2)	For demonstration purposes only, you can enable monitor-only mode for the service policy, which forwards a copy of traffic to the ASA CX module, while the original traffic remains unaffected. Another option for demonstration purposes is to configure a traffic-forwarding interface instead of a service policy in monitor-only mode. The traffic-forwarding interface sends all traffic directly to the ASA CX module, bypassing the ASA. We modified the following screen: Configuration > Firewall > Service Policy Rules > Add Service Policy Rule > Rule Actions > ASA CX Inspection. The traffic-forwarding feature is supported by CLI only.
NAT 64 support for the ASA CX module	ASA 9.1(2) ASA CX 9.1(2)	You can now use NAT 64 in conjunction with the ASA CX module. We did not modify any screens.
ASA 5585-X with SSP-40 and -60 support for the ASA CX SSP-40 and -60	ASA 9.1(3) ASA CX 9.2(1)	ASA CX SSP-40 and -60 modules can be used with the matching level ASA 5585-X with SSP-40 and -60. We did not modify any screens.
Multiple context mode support for the ASA CX module	ASA 9.1(3) ASA CX 9.2(1)	You can now configure ASA CX service policies per context on the ASA. **Note** Although you can configure per context ASA service policies, the ASA CX module itself (configured in PRSM) is a single context mode device; the context-specific traffic coming from the ASA is checked against the common ASA CX policy. We did not modify any screens.
Filtering packets captured on the ASA CX backplane	ASA 9.1(3) ASA CX 9.2(1)	You can now filter packets captured on the ASA CX backplane using the **match** or **access-list** keyword with the **capture interface asa_dataplane** command. Control traffic specific to the ASA CX module is not affected by the access-list or match filtering; the ASA captures all control traffic. In multiple context mode, configure the packet capture per context. Note that all control traffic in multiple context mode goes only to the system execution space. Because control traffic cannot be filtered using an access-list or match, these options are not available in the system execution space. We did not modify any ASDM screens.

CHAPTER **28**

ASA IPS Module

This chapter describes how to configure the ASA IPS module. The ASA IPS module might be a hardware module or a software module, depending on your ASA model. For a list of supported ASA IPS modules per ASA model, see the *Cisco ASA Compatibility Matrix*:

http://www.cisco.com/en/US/docs/security/asa/compatibility/asamatrx.html

This chapter includes the following sections:

- Information About the ASA IPS Module, page 28-1
- Licensing Requirements for the ASA IPS module, page 28-5
- Guidelines and Limitations, page 28-5
- Default Settings, page 28-6
- Configuring the ASA IPS module, page 28-7
- Managing the ASA IPS module, page 28-19
- Monitoring the ASA IPS module, page 28-24
- Feature History for the ASA IPS module, page 28-25

Information About the ASA IPS Module

The ASA IPS module runs advanced IPS software that provides proactive, full-featured intrusion prevention services to stop malicious traffic, including worms and network viruses, before they can affect your network. This section includes the following topics:

- How the ASA IPS Module Works with the ASA, page 28-2
- Operating Modes, page 28-2
- Using Virtual Sensors (ASA 5512-X and Higher), page 28-3
- Information About Management Access, page 28-4

How the ASA IPS Module Works with the ASA

The ASA IPS module runs a separate application from the ASA. The ASA IPS module might include an external management interface so you can connect to the ASA IPS module directly; if it does not have a management interface, you can connect to the ASA IPS module through the ASA interface. The ASA IPS SSP on the ASA 5585-X includes data interfaces; these interfaces provide additional port-density for the ASA. However, the overall through-put of the ASA is not increased.

Traffic goes through the firewall checks before being forwarded to the ASA IPS module. When you identify traffic for IPS inspection on the ASA, traffic flows through the ASA and the ASA IPS module as follows. **Note**: This example is for "inline mode." See Operating Modes, page 28-2 for information about "promiscuous mode," where the ASA only sends a copy of the traffic to the ASA IPS module.

1. Traffic enters the ASA.

2. Incoming VPN traffic is decrypted.

3. Firewall policies are applied.

4. Traffic is sent to the ASA IPS module.

5. The ASA IPS module applies its security policy to the traffic, and takes appropriate actions.

6. Valid traffic is sent back to the ASA; the ASA IPS module might block some traffic according to its security policy, and that traffic is not passed on.

7. Outgoing VPN traffic is encrypted.

8. Traffic exits the ASA.

Figure 28-1 shows the traffic flow when running the ASA IPS module in inline mode. In this example, the ASA IPS module automatically blocks traffic that it identified as an attack. All other traffic is forwarded through the ASA.

Figure 28-1   ASA IPS module Traffic Flow in the ASA: Inline Mode
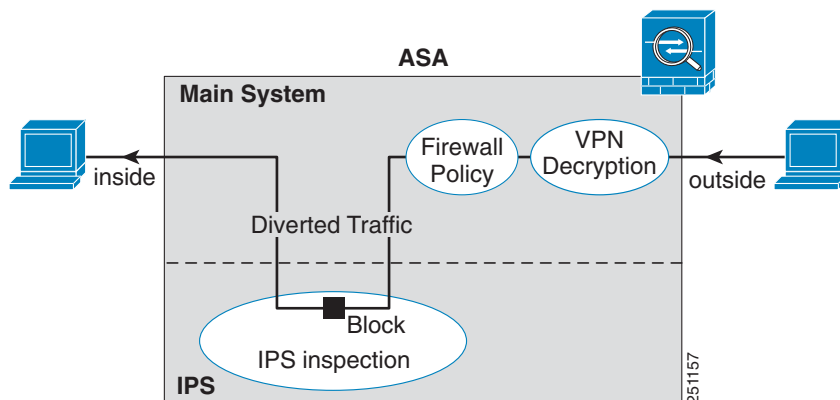
Operating Modes

You can send traffic to the ASA IPS module using one of the following modes:

- Inline mode—This mode places the ASA IPS module directly in the traffic flow (see Figure 28-1). No traffic that you identified for IPS inspection can continue through the ASA without first passing through, and being inspected by, the ASA IPS module. This mode is the most secure because every

packet that you identify for inspection is analyzed before being allowed through. Also, the ASA IPS module can implement a blocking policy on a packet-by-packet basis. This mode, however, can affect throughput.

- Promiscuous mode—This mode sends a duplicate stream of traffic to the ASA IPS module. This mode is less secure, but has little impact on traffic throughput. Unlike inline mode, in promiscuous mode the ASA IPS module can only block traffic by instructing the ASA to shun the traffic or by resetting a connection on the ASA. Also, while the ASA IPS module is analyzing the traffic, a small amount of traffic might pass through the ASA before the ASA IPS module can shun it. Figure 28-2 shows the ASA IPS module in promiscuous mode. In this example, the ASA IPS module sends a shun message to the ASA for traffic it identified as a threat.

Figure 28-2   ASA IPS module Traffic Flow in the ASA: Promiscuous Mode

Using Virtual Sensors (ASA 5512-X and Higher)

The ASA IPS module running IPS software Version 6.0 and later can run multiple virtual sensors, which means you can configure multiple security policies on the ASA IPS module. You can assign each ASA security context or single mode ASA to one or more virtual sensors, or you can assign multiple security contexts to the same virtual sensor. See the IPS documentation for more information about virtual sensors, including the maximum number of sensors supported.

Figure 28-3 shows one security context paired with one virtual sensor (in inline mode), while two security contexts share the same virtual sensor.

Figure 28-3   Security Contexts and Virtual Sensors
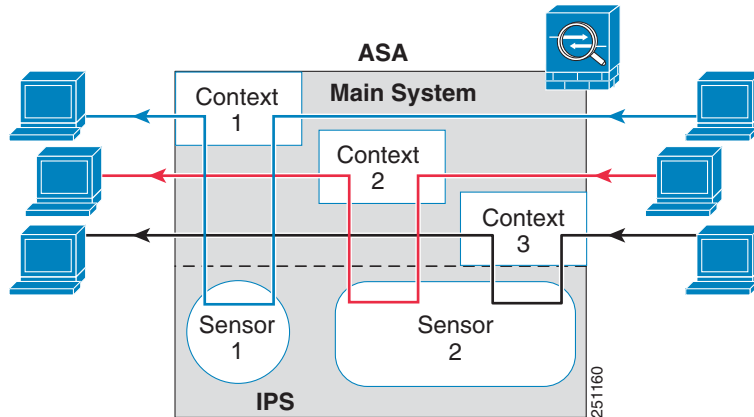
Figure 28-4 shows a single mode ASA paired with multiple virtual sensors (in inline mode); each defined traffic flow goes to a different sensor.

Figure 28-4   Single Mode ASA with Multiple Virtual Sensors
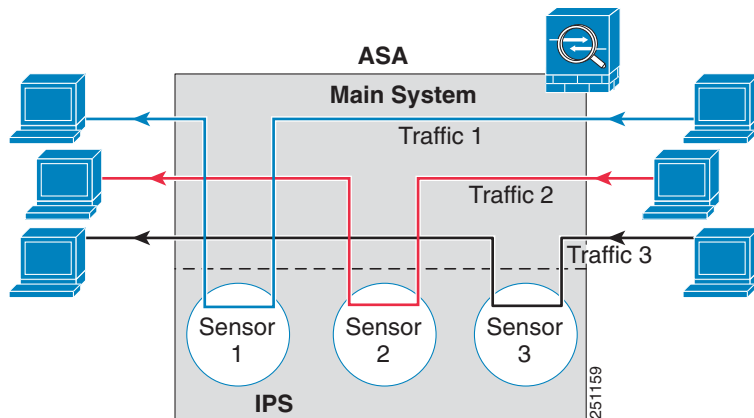
Information About Management Access

You can manage the IPS application using the following methods:

- Sessioning to the module from the ASA—If you have CLI access to the ASA, then you can session to the module and access the module CLI. See Sessioning to the Module from the ASA (May Be Required), page 28-11.

- Connecting to the IPS management interface using ASDM or SSH—After you launch ASDM from the ASA, your management station connects to the module management interface to configure the IPS application. For SSH, you can access the module CLI directly on the module management interface. (Telnet access requires additional configuration in the module application). The module management interface can also be used for sending syslog messages or allowing updates for the module application, such as signature database updates.

See the following information about the management interface:

- ASA 5585-X—The IPS management interface is a separate external Gigabit Ethernet interface.

– ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, ASA 5555-X—These models run the ASA IPS module as a software module. The IPS management interface shares the Management 0/0 interface with the ASA. Separate MAC addresses and IP addresses are supported for the ASA and ASA IPS module. You must perform configuration of the IPS IP address within the IPS operating system (using the CLI or ASDM). However, physical characteristics (such as enabling the interface) are configured on the ASA. You can remove the ASA interface configuration (specifically the interface name) to dedicate this interface as an IPS-only interface. This interface is management-only.

– ASA 5505—You can use an ASA VLAN to allow access to an internal management IP address over the backplane.

Licensing Requirements for the ASA IPS module

The following table shows the licensing requirements for this feature:

Model	License Requirement
ASA 5505	Base License.
ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, ASA 5555-X	IPS Module License. **Note** The IPS module license lets you run the IPS software module on the ASA. You must also purchase a separate IPS signature subscription; for failover, purchase a subscription for each unit. To obtain IPS signature support, you must purchase the ASA with IPS pre-installed (the part number must include "IPS"). The combined failover cluster license does not let you pair non-IPS and IPS units. For example, if you buy the IPS version of the ASA 5515-X (part number ASA5515-IPS-K9) and try to make a failover pair with a non-IPS version (part number ASA5515-K9), then you will not be able to obtain IPS signature updates for the ASA5515-K9 unit, even though it has an IPS module license inherited from the other unit.
ASA 5585-X	Base License.
All other models	No support.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

The ASA 5505 does not support multiple context mode, so multiple context features, such as virtual sensors, are not supported on the AIP SSC.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Model Guidelines

• See the *Cisco ASA Compatibility Matrix* for information about which models support which modules:

http://www.cisco.com/en/US/docs/security/asa/compatibility/asamatrx.html

- The ASA 5505 does not support multiple context mode, so multiple context features, such as virtual sensors, are not supported on the AIP SSC.

- The ASA IPS module for the ASA 5512-X and higher supports higher performance requirements, while the ASA IPS module for the ASA 5505 is designed for a small office installation. The following features are not supported for the ASA 5505:

  – Virtual sensors

  – Anomaly detection

  – Unretirement of default retired signatures

Additional Guidelines

- The total throughput for the ASA plus the IPS module is lower than ASA throughput alone.

  – ASA 5512-X through ASA 5555-X—See http://www.cisco.com/en/US/prod/collateral/vpndevc/ps6032/ps6094/ps6120/qa_c67-700608.html

  – ASA 5585-X—See http://www.cisco.com/en/US/prod/collateral/vpndevc/ps6032/ps6094/ps6120/qa_c67-617018.html

  – ASA 5505—See http://www.cisco.com/en/US/prod/collateral/vpndevc/ps6032/ps6094/ps6120/product_data_sheet0900aecd802930c5.html

- You cannot change the software type installed on the module; if you purchase an ASA IPS module, you cannot later install other software on it.

Default Settings

Table 28-1 lists the default settings for the ASA IPS module.

Table 28-1   Default Network Parameters

Parameters	Default
Management VLAN (ASA 5505 only)	VLAN 1
Management IP address	192.168.1.2/24
Gateway	192.168.1.1/24 (the default ASA management IP address)
Username	cisco
Password	cisco

Note The default management IP address on the ASA is 192.168.1.1/24.

Configuring the ASA IPS module

This section describes how to configure the ASA IPS module and includes the following topics:

- Task Flow for the ASA IPS Module, page 28-7
- Connecting the ASA IPS Management Interface, page 28-8
- Sessioning to the Module from the ASA (May Be Required), page 28-11
- Configuring Basic IPS Module Network Settings, page 28-12
- (ASA 5512-X through ASA 5555-X) Booting the Software Module, page 28-12
- Configuring the Security Policy on the ASA IPS Module, page 28-15
- Assigning Virtual Sensors to a Security Context (ASA 5512-X and Higher), page 28-17
- Diverting Traffic to the ASA IPS module, page 28-18

Task Flow for the ASA IPS Module

Configuring the ASA IPS module is a process that includes configuration of the IPS security policy on the ASA IPS module and then configuration of the ASA to send traffic to the ASA IPS module. To configure the ASA IPS module, perform the following steps:

Step 1 Cable the ASA IPS management interface. See Connecting the ASA IPS Management Interface, page 28-8.

Step 2 Session to the module. Access the IPS CLI over the backplane. For ASDM users, you may need to session to the module to boot the IPS software if it is not running. See Sessioning to the Module from the ASA (May Be Required), page 28-11.

Step 3 (ASA 5512-X through ASA 5555-X; may be required) Install the software module. See (ASA 5512-X through ASA 5555-X) Booting the Software Module, page 28-12.

Step 4 Depending on your ASA model:
- (ASA 5512-X and higher) Configure basic network settings for the IPS module. See (ASA 5512-X and Higher) Configuring Basic Network Settings, page 28-13.
- (ASA 5505) Configure the management VLAN and IP address for the IPS module. See (ASA 5505) Configuring Basic Network Settings, page 28-14.

Step 5 On the module, configure the inspection and protection policy, which determines how to inspect traffic and what to do when an intrusion is detected. See Configuring the Security Policy on the ASA IPS Module, page 28-15.

Step 6 (ASA 5512-X and higher, optional) On the ASA in multiple context mode, specify which IPS virtual sensors are available for each context (if you configured virtual sensors). See Assigning Virtual Sensors to a Security Context (ASA 5512-X and Higher), page 28-17.

Step 7 On the ASA, identify traffic to divert to the ASA IPS module. See Diverting Traffic to the ASA IPS module, page 28-18.

Connecting the ASA IPS Management Interface

In addition to providing management access to the IPS module, the IPS management interface needs access to an HTTP proxy server or a DNS server and the Internet so it can download global correlation, signature updates, and license requests. This section describes recommended network configurations. Your network may differ.

- ASA 5585-X (Hardware Module), page 28-8
- ASA 5512-X through ASA 5555-X (Software Module), page 28-9
- ASA 5505, page 28-10

ASA 5585-X (Hardware Module)

The IPS module includes a separate management interface from the ASA.

If you have an inside router

If you have an inside router, you can route between the management network, which can include both the ASA Management 0/0 and IPS Management 1/0 interfaces, and the ASA inside network. Be sure to also add a route on the ASA to reach the Management network through the inside router.

If you do not have an inside router

If you have only one inside network, then you cannot also have a separate management network, which would require an inside router to route between the networks. In this case, you can manage the ASA from the inside interface instead of the Management 0/0 interface. Because the IPS module is a separate device from the ASA, you can configure the IPS Management 1/0 address to be on the same network as the inside interface.
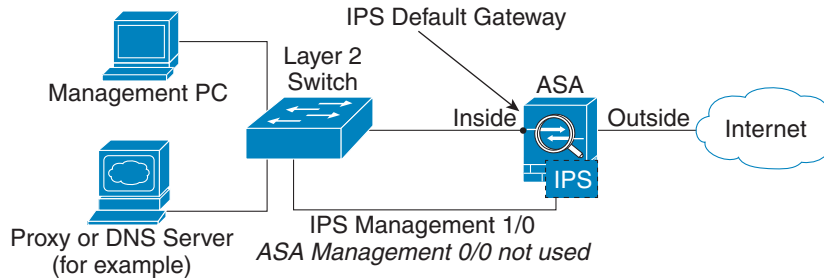
ASA 5512-X through ASA 5555-X (Software Module)

These models run the IPS module as a software module, and the IPS management interface shares the Management 0/0 interface with the ASA.

If you have an inside router

If you have an inside router, you can route between the Management 0/0 network, which includes both the ASA and IPS management IP addresses, and the inside network. Be sure to also add a route on the ASA to reach the Management network through the inside router.

■
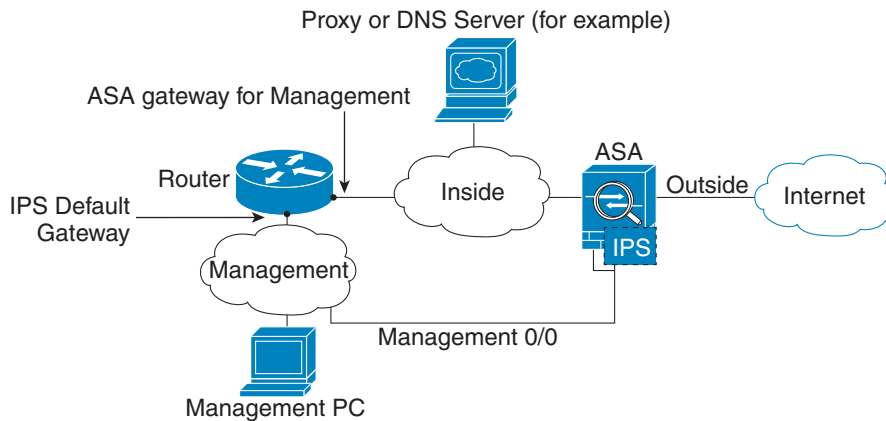
If you do not have an inside router

If you have only one inside network, then you cannot also have a separate management network. In this case, you can manage the ASA from the inside interface instead of the Management 0/0 interface. If you remove the ASA-configured name from the Management 0/0 interface, you can still configure the IPS IP address for that interface. Because the IPS module is essentially a separate device from the ASA, you *can* configure the IPS management address to be on the same network as the inside interface.

> **Note**   You must remove the ASA-configured name for Management 0/0; if it is configured on the ASA, then the IPS address must be on the same network as the ASA, and that excludes any networks already configured on other ASA interfaces. If the name is not configured, then the IPS address can be on any network, for example, the ASA inside network.

ASA 5505

The ASA 5505 does not have a dedicated management interface. You must use an ASA VLAN to access an internal management IP address over the backplane. Connect the management PC to one of the following ports: Ethernet 0/1 through 0/7, which are assigned to VLAN 1.

What to Do Next

- (ASA 5512-X and higher) Configure basic network settings. See (ASA 5512-X and Higher) Configuring Basic Network Settings, page 28-13.

- (ASA 5505) Configure management interface settings. See (ASA 5505) Configuring Basic Network Settings, page 28-14.

Sessioning to the Module from the ASA (May Be Required)

To access the IPS module CLI from the ASA, you can session from the ASA. For software modules, you can either session to the module (using Telnet) or create a virtual console session. A console session might be useful if the control plane is down and you cannot establish a Telnet session.

You may need to access the CLI if you are using multiple context mode and you need to set basic network settings using the CLI, or for troubleshooting.

Detailed Steps

Command	Purpose

Telnet session.
 	
 For a hardware module (for example, the ASA 5585-X):
 	
 session 1
 	
 For a software module (for example, the ASA 5545-X):
 	
 session ips
 	
 Example:
 `hostname# session 1`
 	
 `Opening command session with slot 1.`
 `Connected to slot 1. Escape character`
 `sequence is 'CTRL-^X'.`
 	
 `sensor login: cisco`
 `Password: cisco`
	Accesses the module using Telnet. You are prompted for the username and password. The default username is **cisco**, and the default password is **cisco**.
	
	 Note The first time you log in to the module, you are prompted to change the default password. Passwords must be at least eight characters long and cannot be a word in the dictionary.

Console session (software module only).
 	
 session ips console
 	
 Example:
 `hostname# session ips console`
 	
 `Establishing console session with slot 1`
 `Opening console session with module ips.`
 `Connected to module ips. Escape character`
 `sequence is 'CTRL-SHIFT-6 then x'.`
 	
 `sensor login: cisco`
 `Password: cisco`
	Accesses the module console. You are prompted for the username and password. The default username is **cisco**, and the default password is **cisco**.
	
	 Note Do not use this command in conjunction with a terminal server where **Ctrl-Shift-6, x** is the escape sequence to return to the terminal server prompt. **Ctrl-Shift-6, x** is also the sequence to escape the IPS console and return to the ASA prompt. Therefore, if you try to exit the IPS console in this situation, you instead exit all the way to the terminal server prompt. If you reconnect the terminal server to the ASA, the IPS console session is still active; you can never exit to the ASA prompt. You must use a direct serial connection to return the console to the ASA prompt.
	
	 Use the **session ips** command instead.

(ASA 5512-X through ASA 5555-X) Booting the Software Module

Your ASA typically ships with IPS module software present on Disk0. If the module is not running, or if you are adding the IPS module to an existing ASA, you must boot the module software. If you are unsure if the module is running, you will not see the IPS Basic Configuration screen when you run the Startup Wizard (see Configuring Basic IPS Module Network Settings, page 28-12).

Detailed Steps

Step 1 Do one of the following:

- New ASA with IPS pre-installed—To view the IPS module software filename in flash memory, choose **Tools > File Management**.

  For example, look for a filename like IPS-SSP_5512-K9-sys-1.1-a-7.1-4-E4.aip. Note the filename; you will need this filename later in the procedure.

- Existing ASA with new IPS installation—Download the IPS software from Cisco.com to your computer. If you have a Cisco.com login, you can obtain the software from the following website:

  http://www.cisco.com/cisco/software/navigator.html?mdfid=282164240

  Choose **Tools > File Management**, then choose **File Transfer > Between Local PC and Flash** to upload the new image to disk0. Note the filename; you will need this filename later in the procedure.

Step 2 Choose **Tools > Command Line Interface**.

Step 3 To set the IPS module software location in disk0, enter the following command and then click **Send**:

```
sw-module module ips recover configure image disk0:file_path
```

For example, using the filename in the example in Step 1, enter:

```
sw-module module ips recover configure image disk0:IPS-SSP_5512-K9-sys-1.1-a-7.1-4-E4.aip
```

Step 4 To install and load the IPS module software, enter the following command and then click **Send**:

```
sw-module module ips recover boot
```

Step 5 To check the progress of the image transfer and module restart process, enter the following command and then click **Send**:

```
show module ips details
```

The Status field in the output indicates the operational status of the module. A module operating normally shows a status of "Up." While the ASA transfers an application image to the module, the Status field in the output reads "Recover." When the ASA completes the image transfer and restarts the module, the newly transferred image is running.

Configuring Basic IPS Module Network Settings

(ASA 5512-X and Higher) Configuring Basic Network Settings

In single context mode, you can use the Startup Wizard in ASDM to configure basic IPS network configuration. These settings are saved to the IPS configuration, not the ASA configuration.

In multiple context mode, session to the module from the ASA and configure basic settings using the **setup** command.

> **Note**   (ASA 5512-X through ASA 5555-X) If you do not see the IPS Basic Configuration screen in your wizard, then the IPS module is not running. See (ASA 5512-X through ASA 5555-X) Booting the Software Module, page 28-12, and then repeat this procedure after you install the module.

Detailed Steps—Single Mode

Step 1   Choose **Wizards > Startup Wizard**.

Step 2   Click **Next** to advance through the initial screens until you reach the IPS Basic Configuration screen.

Step 3   In the Network Settings area, configure the following:

- IP Address—The management IP address. By default, the address is 192.168.1.2.
- Subnet Mask—The subnet mask for the management IP address.
- Gateway—The IP address of the upstream router. The IP address of the next hop router. See Connecting the ASA IPS Management Interface, page 28-8 to understand the requirements for your network. The default setting of the ASA management IP address will not work.
- HTTP Proxy Server—(Optional) The HTTP proxy server address. You can use a proxy server to download global correlation updates and other information instead of downloading over the Internet.
- HTTP Proxy Port—(Optional) The HTTP proxy server port.
- DNS Primary—(Optional) The primary DNS server address. If you are using a DNS server, you must configure at least one DNS server and it must be reachable for global correlation updates to be successful.

  For global correlation to function, you must have either a DNS server or an HTTP proxy server configured at all times. DNS resolution is supported only for accessing the global correlation update server.

Step 4   In the Management Access List area, enter an IP address and subnet mask for any hosts that are allowed to access the IPS management interface, and click **Add**. You can add multiple IP addresses.

Step 5   In the Cisco Account Password area, set the password for the username **cisco** and confirm it. The username **cisco** and this password are used for Telnet sessions from hosts specified by the management ACL and when accessing the IPS module from ASDM (Configuration > IPS). By default, the password is **cisco**.

Step 6   In the Network Participation area, which you use to have the IPS module participate in SensorBase data sharing, click **Full**, **Partial**, or **Off**.

Detailed Steps—Multiple Mode Using the CLI

	Command	Purpose
Step 1	Session to the IPS module according to the Sessioning to the Module from the ASA (May Be Required), page 28-11.	
Step 2	**setup** **Example:** sensor# setup	Runs the setup utility for initial configuration of the ASA IPS module. You are prompted for basic settings. For the default gateway, specify the IP address of the upstream router. See Connecting the ASA IPS Management Interface, page 28-8 to understand the requirements for your network. The default setting of the ASA management IP address will not work.

(ASA 5505) Configuring Basic Network Settings

An ASA IPS module on the ASA 5505 does not have any external interfaces. You can configure a VLAN to allow access to an internal IPS management IP address over the backplane. By default, VLAN 1 is enabled for IPS management. You can only assign one VLAN as the management VLAN. This section describes how to change the management VLAN and IP address if you do not want to use the default, and how to set other required network parameters.

Note   Perform this configuration on the ASA 5505, not on the ASA IPS module.

Prerequisites

When you change the IPS VLAN and management address from the default, be sure to also configure the matching ASA VLAN and switch port(s) according to the procedures listed in the general operations configuration guide. You must define and configure the VLAN for the ASA so the IPS management interface is accessible on the network.

Restrictions

Do not configure NAT for the management address if you intend to access it using ASDM. For initial setup with ASDM, you need to access the real address. After initial setup (where you set the password on the ASA IPS module), you can configure NAT and supply ASDM with the translated address for accessing the ASA IPS module.

Detailed Steps

Step 1   In ASDM, choose **Configuration > Device Setup > SSC Setup**.

Note   The following settings are written to the ASA IPS module application configuration, not the ASA configuration.

Step 2   In the Management Interface area, set the following:

a. Choose the Interface VLAN from the drop-down list.

This setting allows you to manage the ASA IPS module using this VLAN.

    b. Enter the IPS management IP address. Make sure this address is on the same subnet as the ASA VLAN IP address. For example, if you assigned 10.1.1.1 to the VLAN for the ASA, then assign another address on that network, such as 10.1.1.2, for the IPS management address. By default, the address is 192.168.1.2

    c. Choose the subnet mask from the drop-down list.

    d. Enter the default gateway IP address.

    Set the gateway to be the ASA IP address for the management VLAN. By default, this IP address is 192.168.1.1.

Step 3 In the Management Access List area, enter the following:

    a. Enter the IP address for the management host network.

    b. Choose the subnet mask from the drop-down list.

    c. Click **Add** to add these settings to the Allowed Hosts/Networks list.

Step 4 In the IPS Password area, do the following:

    a. Enter the current password. The default password is **cisco**.

    b. Enter the new password, and confirm the change.

Step 5 Click **Apply** to save the settings to the running configuration.

Step 6 To launch the IPS Startup Wizard, click the **Configure the IPS SSC module** link.

Configuring the Security Policy on the ASA IPS Module

This section describes how to configure the ASA IPS module application.

Detailed Steps

Step 1 Connect to ASDM using the *ASA* management IP address.

Step 2 To access the IPS Device Manager (IDM) from ASDM, click **Configuration > IPS**.

Step 3   Enter the IP address, username and password that you set in the Configuring Basic IPS Module Network Settings, page 28-12, as well as the port. The default IP address and port is 192.168.1.2:443. The default username and password is **cisco** and **cisco**.

If the password to access IDM is lost, you can reset the password using ASDM. See Resetting the Password, page 28-23, for more information.

Step 4   To save the login information on your local PC, check the **Save IPS login information on local host** check box.

Step 5   Click **Continue**.

The Startup Wizard pane appears.

Step 6   Click **Launch Startup Wizard**. Complete the screens as prompted. For more information, see the IDM online help.
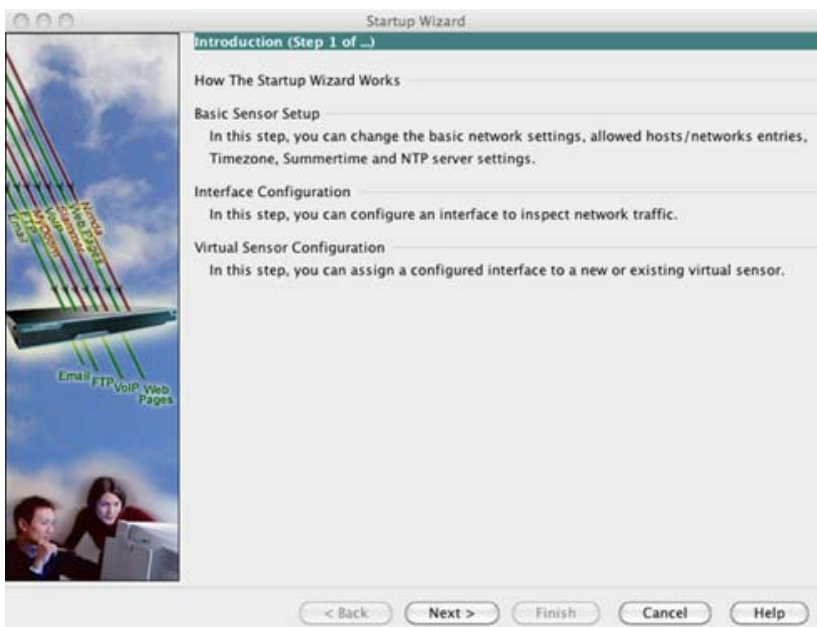
(ASA 5512-X and higher) If you configure virtual sensors, you identify one of the sensors as the default. If the ASA series does not specify a virtual sensor name in its configuration, the default sensor is used.

What to Do Next

- For the ASA in multiple context mode, see Assigning Virtual Sensors to a Security Context (ASA 5512-X and Higher), page 28-17.

- For the ASA in single context mode, see Diverting Traffic to the ASA IPS module, page 28-18.

Assigning Virtual Sensors to a Security Context (ASA 5512-X and Higher)

If the ASA is in multiple context mode, then you can assign one or more IPS virtual sensors to each context. Then, when you configure the context to send traffic to the ASA IPS module, you can specify a sensor that is assigned to the context; you cannot specify a sensor that you did not assign to the context. If you do not assign any sensors to a context, then the default sensor configured on the ASA IPS module is used. You can assign the same sensor to multiple contexts.

> **Note** You do not need to be in multiple context mode to use virtual sensors; you can be in single mode and use different sensors for different traffic flows.

Prerequisites

For more information about configuring contexts, see the general operations configuration guide.

Detailed Steps

Step 1 In the ASDM Device List pane, double-click **System** under the active device IP address.

Step 2 On the Context Management > Security Contexts pane, choose a context that you want to configure, and click **Edit**.

The Edit Context dialog box appears. For more information about configuring contexts, see the general operations configuration guide.

Step 3 In the IPS Sensor Allocation area, click **Add**.

The IPS Sensor Selection dialog box appears.

Step 4 From the Sensor Name drop-down list, choose a sensor name from those configured on the ASA IPS module.

Step 5 (Optional) To assign a mapped name to the sensor, enter a value in the Mapped Sensor Name field.

This sensor name can be used within the context instead of the actual sensor name. If you do not specify a mapped name, the sensor name is used within the context. For security purposes, you might not want the context administrator to know which sensors are being used by the context. Or you might want to genericize the context configuration. For example, if you want all contexts to use sensors called "sensor1" and "sensor2," then you can map the "highsec" and "lowsec" sensors to sensor1 and sensor2 in context A, but map the "medsec" and "lowsec" sensors to sensor1 and sensor2 in context B.

Step 6 Click **OK** to return to the Edit Context dialog box.

Step 7 (Optional) To set one sensor as the default sensor for this context, from the Default Sensor drop-down list, choose a sensor name.

If you do not specify a sensor name when you configure IPS within the context configuration, the context uses this default sensor. You can only configure one default sensor per context. If you do not specify a sensor as the default, and the context configuration does not include a sensor name, then traffic uses the default sensor on the ASA IPS module.

Step 8   Repeat this procedure for each security context.

Step 9   Change to each context to configure the IPS security policy as described in Diverting Traffic to the ASA IPS module, page 28-18.

What to Do Next

Change to each context to configure the IPS security policy as described in Diverting Traffic to the ASA IPS module, page 28-18.

Diverting Traffic to the ASA IPS module

This section identifies traffic to divert from the ASA to the ASA IPS module.

Prerequisites

In multiple context mode, perform these steps in each context execution space. To change to a context, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Detailed Steps

Step 1   Choose **Configuration > Firewall > Service Policy Rules**.

Step 2   Choose **Add > Add Service Policy Rule**. The Add Service Policy Rule Wizard - Service Policy dialog box appears.

Step 3   Complete the Service Policy dialog box as desired. See the ASDM online help for more information about these screens.

Step 4   Click **Next**. The Add Service Policy Rule Wizard - Traffic Classification Criteria dialog box appears.

Step 5   Complete the Traffic Classification Criteria dialog box as desired. See the ASDM online help for more information about these screens.

Step 6   Click **Next** to show the Add Service Policy Rule Wizard - Rule Actions dialog box.

Step 7   Click the **Intrusion Prevention** tab.

Step 8   Check the **Enable IPS for this traffic flow** check box.

Step 9   In the Mode area, click **Inline Mode** or **Promiscuous Mode**. See Operating Modes, page 28-2 for more information.

Step 10   In the If IPS Card Fails area, click **Permit traffic** or **Close traffic**. The Close traffic option sets the ASA to block all traffic if the ASA IPS module is unavailable. The Permit traffic option sets the ASA to allow all traffic through, uninspected, if the ASA IPS module is unavailable. For information about the IPS Sensor Selection area, see the ASDM online help.

Step 11   (ASA 5512-X and higher) From the IPS Sensor to use drop-down list, choose a virtual sensor name.

If you use virtual sensors, you can specify a sensor name using this option. If you use multiple context mode on the ASA, you can only specify sensors that you assigned to the context (see Assigning Virtual Sensors to a Security Context (ASA 5512-X and Higher), page 28-17). If you do not specify a sensor name, then the traffic uses the default sensor. In multiple context mode, you can specify a default sensor for the context. In single mode or if you do not specify a default sensor in multiple mode, the traffic uses the default sensor that is set on the ASA IPS module.

Step 12   Click **OK** and then **Apply**.

Step 13   Repeat this procedure to configure additional traffic flows as desired.

Managing the ASA IPS module

This section includes procedures that help you recover or troubleshoot the module and includes the following topics:

Installing and Booting an Image on the Module

If the module suffers a failure, and the module application image cannot run, you can reinstall a new image on the module from a TFTP server (for a hardware module), or from the local disk (software module).

Note   Do not use the **upgrade** command within the module software to install the image.

Prerequisites

- Hardware module—Be sure the TFTP server that you specify can transfer files up to 60 MB in size.

  Note   This process can take approximately 15 minutes to complete, depending on your network and the size of the image.

- Software module—Copy the image to the ASA internal flash (disk0) before completing this procedure.

  Note   Before you download the IPS software to disk0, make sure at least 50% of the flash memory is free. When you install IPS, IPS reserves 50% of the internal flash memory for its file system.

Detailed Steps

	Command	Purpose
Step 1	For a hardware module (for example, the ASA 5585-X): **hw-module module 1 recover configure** For a software module (for example, the ASA 5545-X): **sw-module module ips recover configure image disk0:**_file_path_ **Example:** hostname# hw-module module 1 recover configure Image URL [tftp://127.0.0.1/myimage]: tftp://10.1.1.1/ids-newimg Port IP Address [127.0.0.2]: 10.1.2.10 Port Mask [255.255.255.254]: 255.255.255.0 Gateway IP Address [1.1.2.10]: 10.1.2.254 VLAN ID [0]: 100	Specifies the location of the new image. For a hardware module—This command prompts you for the URL for the TFTP server, the management interface IP address and netmask, gateway address, and VLAN ID (ASA 5505 only). These network parameters are configured in ROMMON; the network parameters you configured in the module application configuration are not available to ROMMON, so you must set them separately here. For a software module—Specify the location of the image on the local disk. You can view the recovery configuration using the **show module {1 | ips} recover** command. In multiple context mode, enter this command in the system execution space.
Step 2	For a hardware module: **hw-module module 1 recover boot** For a software module: **sw-module module ips recover boot** **Example:** hostname# hw-module module 1 recover boot	Installs and boots the IPS module software.
Step 3	For a hardware module: **show module 1 details** For a software module: **show module ips details** **Example:** hostname# show module 1 details	Checks the progress of the image transfer and module restart process. The Status field in the output indicates the operational status of the module. A module operating normally shows a status of "Up." While the ASA transfers an application image to the module, the Status field in the output reads "Recover." When the ASA completes the image transfer and restarts the module, the newly transferred image is running.

Shutting Down the Module

Shutting down the module software prepares the module to be safely powered off without losing configuration data. **Note**: If you reload the ASA, the module is not automatically shut down, so we recommend shutting down the module before reloading the ASA. To gracefully shut down the module, perform the following steps at the ASA CLI.

Detailed Steps

Command	Purpose

For a hardware module (for example, the ASA 5585-X):
 	
`hw-module module 1 shutdown`
 	
For a software module (for example, the ASA 5545-X):
 	
`sw-module module ips shutdown`
 	
Example:
`hostname# hw-module module 1 shutdown`
	Shuts down the module.

Uninstalling a Software Module Image

To uninstall a software module image and associated configuration, perform the following steps.

Detailed Steps

	Command	Purpose

Step 1	`sw-module module ips uninstall`
	
	Example:
	`hostname# sw-module module ips uninstall`
	`Module ips will be uninstalled. This will`
	`completely remove the`
	`disk image associated with the sw-module`
	`including any configuration`
	`that existed within it.`
	
	`Uninstall module <id>? [confirm]`
		Permanently uninstalls the software module image and associated configuration.

Step 2	`reload`
	
	Example:
	`hostname# reload`
		Reloads the ASA. You must reload the ASA before you can install a new module type.

Resetting the Password

You can reset the module password to the default. For the user **cisco**, the default password is **cisco**. After resetting the password, you should change it to a unique value using the module application.

Resetting the module password causes the module to reboot. Services are not available while the module is rebooting.

If you cannot connect to ASDM with the new password, restart ASDM and try to log in again. If you defined a new password and still have an existing password in ASDM that is different from the new password, clear the password cache by choosing **File > Clear ASDM Password Cache**, then restart ASDM and try to log in again.

To reset the module password to the default of cisco, perform the following steps.

Detailed Steps

Step 1   From the ASDM menu bar, choose **Tools >** *module* **Password Reset**.

The Password Reset confirmation dialog box appears.

Step 2   Click **OK** to reset the password to the default.

A dialog box displays the success or failure of the password reset.

Step 3   Click **Close** to close the dialog box.

Reloading or Resetting the Module

To reload or reset the module, enter one of the following commands at the ASA CLI.

Detailed Steps

Command	Purpose

For a hardware module (for example, the ASA 5585-X):
 	
`hw-module module 1 reload`
 	
For a software module (for example, the ASA 5545-X):
 	
`sw-module module ips reload`
 	
Example:
`hostname# hw-module module 1 reload`
	Reloads the module software.

For a hardware module:
 	
`hw-module module 1 reset`
 	
For a software module:
 	
`sw-module module ips reset`
 	
Example:
`hostname# hw-module module 1 reset`
	Performs a reset, and then reloads the module.

Monitoring the ASA IPS module

See the Intrusion Prevention Tab in the general operations configuration guide.

Feature History for the ASA IPS module

Table 28-2 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 28-2   Feature History for the ASA IPS module

Feature Name	Platform Releases	Feature Information
AIP SSM	7.0(1)	We introduced support for the AIP SSM for the ASA 5510, 5520, and 5540. The following screen was introduced: Configuration > Firewall > Service Policy Rules > Add/Edit Service Policy Rule > Intrusion Prevention.
Virtual sensors (ASA 5510 and higher)	8.0(2)	Virtual sensor support was introduced. Virtual sensors let you configure multiple security policies on the ASA IPS module. The following screen was modified: Context Management > Security Contexts > Edit Context.
AIP SSC for the ASA 5505	8.2(1)	We introduced support for the AIP SSC for the ASA 5505. The following screen was introduced: Configuration > Device Setup > SSC Setup.
Support for the ASA IPS SSP-10, -20, -40, and -60 for the ASA 5585-X	8.2(5)/ 8.4(2)	We introduced support for the ASA IPS SSP-10, -20, -40, and -60 for the ASA 5585-X. You can only install the ASA IPS SSP with a matching-level SSP; for example, SSP-10 and ASA IPS SSP-10. **Note** The ASA 5585-X is not supported in Version 8.3.
Support for Dual SSPs for SSP-40 and SSP-60	8.4(2)	For SSP-40 and SSP-60, you can use two SSPs of the same level in the same chassis. Mixed-level SSPs are not supported (for example, an SSP-40 with an SSP-60 is not supported). Each SSP acts as an independent device, with separate configurations and management. You can use the two SSPs as a failover pair if desired. **Note** When using two SSPs in the chassis, VPN is not supported; note, however, that VPN has not been disabled. We did not modify any screens.
Support for the ASA IPS SSP for the ASA 5512-X through ASA 5555-X	8.6(1)	We introduced support for the ASA IPS SSP software module for the ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, and ASA 5555-X. We did not modify any screens.