ASDM User Guide

Version 5.2(2)
November 2006

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Text Part Number: OL-10106-02

ASDM User Guide

INDEX

Preface

The *ASDM User Guide* contains the information that is available in the ASDM online help system.

This preface contains the following topics:

Related Documentation

For more information, refer to the following documentation:

- *Cisco ASDM Release Notes*
- *Cisco Security Appliance Command Line Configuration Guide*
- *Cisco Security Appliance Command Reference*
- *Cisco ASA 5500 Series Adaptive Security Appliance Getting Started Guide*
- *Cisco ASA 5500 Series Release Notes*
- *Cisco Security Appliance Logging Configuration and System Log Messages*

Document Conventions

Command descriptions use these conventions:

- Braces ({ }) indicate a required choice.
- Square brackets ([ ]) indicate optional elements.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- **Boldface** indicates commands and keywords that are entered literally as shown.

- *Italics* indicate arguments for which you supply values.

Examples use these conventions:

- Examples depict screen displays and the command line in `screen` font.
- Information you need to enter in examples is shown in **`boldface screen`** font.
- Variables for which you must supply a value are shown in *`italic screen`* font.

Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.

Obtaining Documentation

Cisco documentation and additional literature are available on Cisco.com. This section explains the product documentation resources that Cisco offers.

Cisco.com

You can access the most current Cisco documentation at this URL:

http://www.cisco.com/techsupport

You can access the Cisco website at this URL:

http://www.cisco.com

You can access international Cisco websites at this URL:

http://www.cisco.com/public/countries_languages.shtml

Product Documentation DVD

The Product Documentation DVD is a library of technical product documentation on a portable medium. The DVD enables you to access installation, configuration, and command guides for Cisco hardware and software products. With the DVD, you have access to the HTML documentation and some of the PDF files found on the Cisco website at this URL:

http://www.cisco.com/univercd/home/home.htm

The Product Documentation DVD is created and released regularly. DVDs are available singly or by subscription. Registered Cisco.com users can order a Product Documentation DVD (product number DOC-DOCDVD= or DOC-DOCDVD=SUB) from Cisco Marketplace at the Product Documentation Store at this URL:

http://www.cisco.com/go/marketplace/docstore

Ordering Documentation

You must be a registered Cisco.com user to access Cisco Marketplace. Registered users may order Cisco documentation at the Product Documentation Store at this URL:

http://www.cisco.com/go/marketplace/docstore

If you do not have a user ID or password, you can register at this URL:

http://tools.cisco.com/RPF/register/register.do

Documentation Feedback

You can provide feedback about Cisco technical documentation on the Cisco Support site area by entering your comments in the feedback form available in every online document.

Cisco Product Security Overview

Cisco provides a free online Security Vulnerability Policy portal at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

From this site, you will find information about how to do the following:

- Report security vulnerabilities in Cisco products
- Obtain assistance with security incidents that involve Cisco products
- Register to receive security information from Cisco

A current list of security advisories, security notices, and security responses for Cisco products is available at this URL:

http://www.cisco.com/go/psirt

To see security advisories, security notices, and security responses as they are updated in real time, you can subscribe to the Product Security Incident Response Team Really Simple Syndication (PSIRT RSS) feed. Information about how to subscribe to the PSIRT RSS feed is found at this URL:

http://www.cisco.com/en/US/products/products_psirt_rss_feed.html

Reporting Security Problems in Cisco Products

Cisco is committed to delivering secure products. We test our products internally before we release them, and we strive to correct all vulnerabilities quickly. If you think that you have identified a vulnerability in a Cisco product, contact PSIRT:

- For emergencies only—security-alert@cisco.com

  An emergency is either a condition in which a system is under active attack or a condition for which a severe and urgent security vulnerability should be reported. All other conditions are considered nonemergencies.

- For nonemergencies—psirt@cisco.com

In an emergency, you can also reach PSIRT by telephone:

- 1 877 228-7302
- 1 408 525-6532

Tip We encourage you to use Pretty Good Privacy (PGP) or a compatible product (for example, GnuPG) to encrypt any sensitive information that you send to Cisco. PSIRT can work with information that has been encrypted with PGP versions 2.*x* through 9.*x*.

Never use a revoked encryption key or an expired encryption key. The correct public key to use in your correspondence with PSIRT is the one linked in the Contact Summary section of the Security Vulnerability Policy page at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

The link on this page has the current PGP key ID in use.

If you do not have or use PGP, contact PSIRT to find other means of encrypting the data before sending any sensitive material.

Product Alerts and Field Notices

Modifications to or updates about Cisco products are announced in Cisco Product Alerts and Cisco Field Notices. You can receive these announcements by using the Product Alert Tool on Cisco.com. This tool enables you to create a profile and choose those products for which you want to receive information.

To access the Product Alert Tool, you must be a registered Cisco.com user. Registered users can access the tool at this URL:

http://tools.cisco.com/Support/PAT/do/ViewMyProfiles.do?local=en

To register as a Cisco.com user, go to this URL:

http://tools.cisco.com/RPF/register/register.do

Obtaining Technical Assistance

Cisco Technical Support provides 24-hour-a-day award-winning technical assistance. The Cisco Support website on Cisco.com features extensive online support resources. In addition, if you have a valid Cisco service contract, Cisco Technical Assistance Center (TAC) engineers provide telephone support. If you do not have a valid Cisco service contract, contact your reseller.

Cisco Support Website

The Cisco Support website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The website is available 24 hours a day at this URL:

http://www.cisco.com/en/US/support/index.html

Access to all tools on the Cisco Support website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a user ID or password, you can register at this URL:

http://tools.cisco.com/RPF/register/register.do

Note Before you submit a request for service online or by phone, use the **Cisco Product Identification Tool** to locate your product serial number. You can access this tool from the Cisco Support website by clicking the **Get Tools & Resources** link, clicking the **All Tools (A-Z)** tab, and then choosing **Cisco Product Identification Tool** from the alphabetical list. This tool offers three search options:

by product ID or model name; by tree view; or, for certain products, by copying and pasting **show** command output. Search results show an illustration of your product with the serial number label location highlighted. Locate the serial number label on your product and record the information before placing a service call.

Tip **Displaying and Searching on Cisco.com**

If you suspect that the browser is not refreshing a web page, force the browser to update the web page by holding down the Ctrl key while pressing **F5**.

To find technical information, narrow your search to look in technical documentation, not the entire Cisco.com website. After using the Search box on the Cisco.com home page, click the **Advanced Search** link next to the Search box on the resulting page and then click the **Technical Support & Documentation** radio button.

To provide feedback about the Cisco.com website or a particular technical document, click **Contacts & Feedback** at the top of any Cisco.com web page.

Submitting a Service Request

Using the online TAC Service Request Tool is the fastest way to open S3 and S4 service requests. (S3 and S4 service requests are those in which your network is minimally impaired or for which you require product information.) After you describe your situation, the TAC Service Request Tool provides recommended solutions. If your issue is not resolved using the recommended resources, your service request is assigned to a Cisco engineer. The TAC Service Request Tool is located at this URL:

http://www.cisco.com/techsupport/servicerequest

For S1 or S2 service requests, or if you do not have Internet access, contact the Cisco TAC by telephone. (S1 or S2 service requests are those in which your production network is down or severely degraded.) Cisco engineers are assigned immediately to S1 and S2 service requests to help keep your business operations running smoothly.

To open a service request by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411
Australia: 1 800 805 227
EMEA: +32 2 704 55 55
USA: 1 800 553 2447

For a complete list of Cisco TAC contacts, go to this URL:

http://www.cisco.com/techsupport/contacts

Definitions of Service Request Severity

To ensure that all service requests are reported in a standard format, Cisco has established severity definitions.

Severity 1 (S1)—An existing network is “down” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Severity 2 (S2)—Operation of an existing network is severely degraded, or significant aspects of your business operations are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Severity 3 (S3)—Operational performance of the network is impaired while most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Severity 4 (S4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- The Cisco Online Subscription Center is the website where you can sign up for a variety of Cisco e-mail newsletters and other communications. Create a profile and then select the subscriptions that you would like to receive. To visit the Cisco Online Subscription Center, go to this URL:

  http://www.cisco.com/offer/subscribe

- The *Cisco Product Quick Reference Guide* is a handy, compact reference tool that includes brief product overviews, key features, sample part numbers, and abbreviated technical specifications for many Cisco products that are sold through channel partners. It is updated twice a year and includes the latest Cisco channel product offerings. To order and find out more about the *Cisco Product Quick Reference Guide*, go to this URL:

  http://www.cisco.com/go/guide

- Cisco Marketplace provides a variety of Cisco books, reference guides, documentation, and logo merchandise. Visit Cisco Marketplace, the company store, at this URL:

  http://www.cisco.com/go/marketplace/

- Cisco Press publishes a wide range of general networking, training, and certification titles. Both new and experienced users will benefit from these publications. For current Cisco Press titles and other information, go to Cisco Press at this URL:

  http://www.ciscopress.com

- *Internet Protocol Journal* is s a quarterly journal published by Cisco for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the *Internet Protocol Journal* at this URL:

  http://www.cisco.com/ipj

- Networking products offered by Cisco, as well as customer support services, can be obtained at this URL:

  http://www.cisco.com/en/US/products/index.html

- Networking Professionals Connection is an interactive website where networking professionals share questions, suggestions, and information about networking products and technologies with Cisco experts and other networking professionals. Join a discussion at this URL:

  http://www.cisco.com/discuss/networking

- "What's New in Cisco Documentation" is an online publication that provides information about the latest documentation releases for Cisco products. Updated monthly, this online publication is organized by product category to direct you quickly to the documentation for your products. You can view the latest release of "What's New in Cisco Documentation" at this URL:

  http://www.cisco.com/univercd/cc/td/doc/abtunicd/136957.htm

- World-class networking training is available from Cisco. You can view current offerings at this URL:

  http://www.cisco.com/en/US/learning/index.html

Welcome to ASDM

Welcome to ASDM, a browser-based, Java applet used to configure and monitor the software on security appliances. ASDM is loaded from the security appliance, then used to configure, monitor, and manage the device.

For more information about this release, see the following topics:

Important Notes

- **CLI Command Support**—With a few exceptions, almost all CLI commands are fully supported by ASDM. For a list of commands ASDM does not support, see **Unsupported Commands**.
- **Multiple ASDM Sessions**—ASDM allows multiple PCs or workstations to each have one browser session open with the same security appliance software. A single security appliance can support up to 5 concurrent ASDM sessions in single, routed mode. Only one session per browser per PC or workstation is supported for a particular security appliance. In multiple context mode, five concurrent ASDM sessions are supported per context, up to a limit of 32 connections total per security appliance.
- **Security Appliance Release**—This release of ASDM requires Version 7.1 and does not run with earlier security appliance releases.
- **Caveats**—Use the Bug Toolkit on cisco.com to view current caveat information. You can access Bug Toolkit at:

  http://www.cisco.com/cgi-bin/Support/Bugtool/launch_bugtool.pl
- **Changing OS Color Schemes**—If you change the color scheme of your operating system while ASDM is running, you should restart ASDM or some ASDM screens might not display correctly.

New in This Release

This section contains the following topics:

For a complete list of new platform and ASDM features, refer to the *Cisco ASDM Release Notes* on Cisco.com.

Features Introduced in the 5.2(1) Release

See the following topics for more information about the new features in the 5.2(1) release:

- Enhanced and new inspection engines. See Service Policy Rules, page 21-1 and Global Objects, page 6-1.
- Sub-second failover and the High Availability and Scalability Wizard. See Failover, page 12-1.
- Packet Tracer tool. See Packet Tracer, page 1-12.
- Traceroute tool. See Traceroute, page 1-15.
- Expanded VPN Support:
  - ZoneLabs Integrity Server. See Zone Labs Integrity Server, page 28-62.
  - Easy VPN Remote. See Easy VPN Remote, page 28-63.
  - Online Certificate Status Protocol (OCSP) support. See Add/Edit Trustpoint Configuration > Revocation Check Tab, page 33-11 and Add/Edit Trustpoint Configuration > OCSP Rules Tab, page 33-13.
- RIP routing enhancements. See RIP, page 14-22.
- Static Route Tracking/Dual ISP support. See Static Routes, page 14-29.
- Web Cache Communication Protocol (WCCP) support. See WCCP, page 25-3.
- ASA 5505 adaptive security appliance Power over Ethernet port support. See Configuring Switch Ports and VLAN Interfaces for the Cisco ASA 5505 Adaptive Security Appliance, page 5-13.

Features Introduced in the 5.2(2) Release

See the following topics for more information about the new features in the 5.2(1) release:

- IDM Integration. See Accessing IDM from ASDM, page 35-1.
- AIP SSM Password Reset. See Resetting the AIP SSM Password, page 35-2.
- CSC SSM Password Reset. See Restoring the Default Password, page 36-12.
- Additional Multicast Feature Support:
  - PIM neighbor-filter. See Neighbor Filter, page 15-13.
  - PIM bidir-neighbor-filter. See Bidirectional Neighbor Filter, page 15-14.
  - PIM old-register-checksum. See the Generate IOS compatible register messages check box in Rendezvous Points, page 15-16.
  - Multicast Boundary. See MBoundary, page 15-8.

  - MFIB forwarding. See PIM bidir-neighbor-filter. See MForwarding, page 15-10.
- Support for HTTP/HTTPS interactive authentication. See Configuring Advanced AAA Features, page 19-12.
- Added UPN (User Principle Name) to the Primary DN Field for tunnel groups. See. Add/Edit Tunnel Group > General Tab > Authorization Tab, page 28-46.
- Per-interface authorization server groups for tunnel groups. See Add/Edit Tunnel Group > General Tab > Authorization Tab, page 28-46.
- Support for Virtual Telnet Server. See Virtual Access, page 11-11.

Unsupported Commands

ASDM supports almost all commands available for the security appliance, but some commands in an existing configuration are ignored by ASDM. Most of these commands can remain in your configuration; see Show Commands Ignored by ASDM on Device for the ignored commands in your configuration.

In the case of the **alias** command, ASDM enters into Monitor-only mode until you remove the command from your configuration.

This section contains the following topics:

- Ignored and View-Only Commands
- Effects of Unsupported Commands
- Other CLI Limitations

Ignored and View-Only Commands

The following table lists commands that ASDM supports in the configuration when added by the CLI, but that cannot be added or edited in ASDM. If ASDM ignores the command, it does not appear in the ASDM GUI at all. If it is view-only, then the command appears in the GUI, but you cannot edit it.

Unsupported Commands	ASDM Behavior
access-list	Ignored if not used.
capture	Ignored
dns-guard	Ignored
established	Ignored.
failover timeout	Ignored.
ipv6, any IPv6 addresses	Ignored.
object-group icmp-type	View-only.
object-group network	Nested group is view-only.
object-group protocol	View-only.
object-group service	Nested group cannot be added.
pager	Ignored.
pim accept-register route-map	Ignored. Only the **list** option can be configured using ASDM

Unsupported Commands	ASDM Behavior
prefix-list	Ignored if not used in an OSPF area.
route-map	Ignored.

service-policy global	Ignored if it uses a **match access-list** class. For example:
	`access-list myacl line 1 extended permit ip any any`
	`class-map mycm`
	`match access-list mycl`
	`policy-map mypm`
	`class mycm`
	`inspect ftp`
	`service-policy mypm global`

sysopt nodnsalias	Ignored.
sysopt uauth allow-http-cache	Ignored.
terminal	Ignored.
virtual	Ignored.

Effects of Unsupported Commands

- If ASDM loads an existing running configuration and finds IPv6-related commands, ASDM displays a dialog box informing you that it does not support IPv6. You cannot configure any IPv6 commands in ASDM, but all other configuration is available.
- If ASDM loads an existing running configuration and finds other unsupported commands, ASDM operation is unaffected. To view the unsupported commands, see Options > Show Commands Ignored by ASDM on Device.
- If ASDM loads an existing running configuration and finds the **alias** command, it enters Monitor-only mode.

  Monitor-only mode allows access to the following functions:

  - The Monitoring area
  - The CLI tool (Tools > Command Line Interface), which lets you use the CLI commands.

  To exit Monitor-only mode, use the CLI tool or access the security appliance console, and remove the **alias** command. You can use outside NAT instead of the **alias** command. See the *Cisco Security Appliance Command Reference* for more information.

Note You might also be in Monitor-only mode because your user account privilege level, indicated in the status bar at the bottom of the main ASDM window, was set up as less than or equal to 3 by your system administrator, which allows Monitor-only mode. For more information, see Configuration > Properties > Device Administration > User Accounts and Configuration > Device Access > AAA Access.

Other CLI Limitations

ASDM does not support discontinuous subnet masks such as 255.255.0.255. For example, you cannot use the following:

```
ip address inside 192.168.2.1 255.255.0.255
```

About the ASDM Window

The ASDM Window is designed to provide easy access to the many features that the security appliance supports. The ASDM Window includes the following:

- Menus—Provides quick access to files, tools, options and help.
- Toolbar—Lets you navigate ASDM. From the toolbar you can access the home page, configuration, and monitoring panels. You can also search for features, save the configuration, get help and navigate back and forth between panels. The Home, Configuration, and Monitoring buttons each open a panel with a variety of useful tools. The home page offers much information at a glance. Configuration and monitoring offer a useful category tree along the left side of the frame, for access to more detailed configuration or monitoring information.
- Status Bar—Shows the time, connection status, user, and privilege level.

Menus

ASDM includes the following menus:

- File Menu
- Options Menu
- Tools Menu
- Wizards Menu
- Help Menu

File Menu

The File menu manages security appliance configurations, and includes the following items:

- Refresh ASDM with the Running Configuration on the Device—Loads a copy of the running configuration to ASDM. Use refresh to make sure ASDM has a current copy of the running configuration.
- Reset Device to the Factory Default Configuration—Restores the configuration to the factory default. See Reset Device to the Factory Default Configuration dialog box for more information.
- Show Running Configuration in New Window—Displays the current running configuration in a new window.
- Save Running Configuration to Flash—Writes a copy of the running configuration to Flash memory.
- Save Running Configuration to TFTP Server—Stores a copy of the current running configuration file on a TFTP server. See the Save Running Configuration to TFTP Server dialog box for more information.
- Save Running Configuration to Standby Unit—Sends a copy of the running configuration file on the primary unit to the running configuration of a failover standby unit.
- Save Internal Log Buffer to Flash—Saves the log buffer to flash memory.

- Print—Prints the current panel. We recommend landscape page orientation when printing rules. If ASDM is running in Netscape Communicator and the user has not yet granted print privileges to the Java applet, a security dialog appears requesting Print privileges. Click **Grant** to grant the applet printing privileges. When using Internet Explorer, permission to print is already granted when you originally accepted the signed applet.
- Clear ASDM Cache—Clears the local ASDM images. ASDM downloads an image locally when you connect to ASDM.
- Clear Internal Log Buffer—Clears the system log message buffer.
- Exit—Exits ASDM.

Reset Device to the Factory Default Configuration

The default configuration includes the minimum commands required to connect to the security appliance using ASDM. This feature is available only for routed firewall mode; transparent mode does not support IP addresses for interfaces, and setting the interface IP address is one of the actions this feature takes. This feature is also only available in single context mode; a security appliance with a cleared configuration does not have any defined contexts to automatically configure using this feature.

This feature clears the current running configuration and then configures several commands. The configured interface depends on your platform. For a platform with a dedicated management interface, the interface is named “management.” For other platforms, the configured interface is Ethernet 1 and named “inside.”

The following commands apply to the dedicated management interface, Management 0/0 (for a platform without a dedicated management interface, the interface is Ethernet 1):

```
interface management 0/0
   ip address 192.168.1.1 255.255.255.0
   nameif management
   security-level 100
   no shutdown
asdm logging informational 100
asdm history enable
http server enable
http 192.168.1.0 255.255.255.0 management
dhcpd address 192.168.1.2-192.168.1.254 management
dhcpd lease 3600
dhcpd ping_timeout 750
dhcpd enable management
```

If you set the IP address in this dialog box, then the **http** command uses the subnet you specify. Similarly, the **dhcpd address** command range consists of addresses within the subnet that you specify.

After you restore the factory default configuration, save it to internal Flash memory using the File > Save Running Configuration to Flash item. This menu item saves the running configuration to the default location for the startup configuration, even if you previously configured the Boot Image/Configuration to set a different location; when the configuration was cleared, this path was also cleared.

Note This command also clears the Add Boot Image configuration, if present, along with the rest of the configuration. The Add Boot Image pane lets you boot from a specific image, including an image on the external Flash memory card. The next time you reload the security appliance after restoring the factory configuration, it boots from the first image in internal Flash memory; if you do not have an image in internal Flash memory, the security appliance does not boot.

Fields

- Use this address for the "*Interface_ID*" interface which will be named as "*name*"—Manually sets the IP address of the management interface, instead of using the default address, 192.168.1.1. For a platform with a dedicated management interface, the interface is named "management." For other platforms, the configured interface is Ethernet 1 and named "inside."
- Management IP Address—Sets the management interface IP address.
- Management subnet mask—Sets the subnet mask of the interface. If you do not set a mask, the security appliance uses the mask appropriate for the IP address class.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Save Running Configuration to TFTP Server

File > Save Running Configuration to TFTP Server > Save Running Configuration to TFTP Server

This dialog box stores a copy of the current running configuration file on a TFTP server.

Fields

- TFTP Server IP Address—Enter the IP address of the TFTP server.
- Configuration File Path—Enter path on the TFTP server where the file will be saved.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Enter Log File Name

File > Save Internal Log Buffer to Flash > Enter Log File Name

Saves the log buffer to flash memory.

Fields

- Use default file name—Saves the log buffer using LOG-YYYY-MM-DD-hhmmss.txt as the file name.
- Use user-specified file name—Saves the log buffer using a file name that you specify.
- Field Name—Enter the file name for the saved log buffer.

Options Menu

The Options menu lets you set ASDM preferences.

- Show Commands Ignored by ASDM on Device—Displays unsupported commands that have been ignored by ASDM. See the Show Commands Ignored by ASDM on Device dialog box for more information.
- Preferences—Changes the behavior of some ASDM functions between sessions using your web browser cookie feature. See the Preferences dialog box for more information.

Show Commands Ignored by ASDM on Device

Options > Show Commands Ignored by ASDM on Device > Show Commands Ignored by ASDM on Device

Some commands are unsupported in ASDM. Typically, they are ignored when encountered by ASDM, and are displayed in the list of unparsed commands invoked by Show Commands Ignored by ASDM on Device.

ASDM does not change or remove these commands from your configuration. See Unsupported Commands for more information.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Preferences

Options > Preferences > Preferences

The Preferences dialog box lets you change the behavior of some ASDM functions between sessions by using your web browser cookie feature.

Fields

- General tab—Sets general preferences.
  - Preview commands before sending to the device check box—Lets you view CLI commands generated by ASDM.
  - Enable Large Fonts (Requires ASDM Restart) check box—Increases the ASDM icon font size, after closing ASDM and reconnecting. Not all fonts are affected.
  - Confirm before exiting from ASDM check box—Displays a prompt when you try to close ASDM to confirm that you want to exit. This option is checked by default.
- Rules Table tab—Sets preferences for the Rules Table.
  - Diplsay settings—Lets you change the way rules are displayed in the Rules Table.
  - Auto expand network and service object groups with specified prefix—Displays the network and service object groups automatically expanded based on the Auto Expand-Prefix.

- Auto Expand-Prefix—Specifies the prefix of the network and service object groups to automatically expand when displayed.
- Show members of network and service object groups—Select to display members of network and service object groups and the group name in the rules table. If the check box is not selected, only the group name is displayed.
- Limit members to—Enter the number of network and service object groups to display. When the object group members are displayed, then display only the first *nn* members.
- Show all actions for service policy rules—Select to display all action in the rules table. When cleared, a summary is displayed.

– Deployment Settings—Lets you configure the behavior the security appliance has when deploying changes to the rules table.

- Issue clear xlate command when deploying access lists—Check to clear the NAT table when deploying a new access lists. This ensures the access lists that are configured on the security appliance are applied to all translated addresses.
- Show filter panel by default—Displays the filter panel by default.
- Show rule diagram panel by default—Displays the rule diagram by default.

- Applications Inspections tab—Sets Application Inspection map options.
  – Prompt to add inspect map before applying changes—Enables a prompt that reminds you the inspection map has not yet been added.
  – Make advanced view the default inspect view—Select to make the advanced view the default application inspection view.
  - Ask to make advanced view the default view—Enables a dialog box that asks to make the advanced view the default application inspection view. Clear to disable the prompt.
- Syslog Color Settings tab—Sets the background and text colors for system log messages displayed on the Home page.
  – Severity column—Lists each severity level.
  – Background Color column—Shows the background color for messages for each severity level. To change the color, click the appropriate row. The Pick a Color dialog box appears.
  – Foreground Color column—Shows the foreground (text) color for messages for each severity level. To change the color, click the appropriate row. The Pick a Color dialog box appears.
  – Restore Default button—Restores the default settings of white background and colored text.

Note Each time a preference is checked or unchecked, the change is written to the .conf file and becomes available for all the other ASDM sessions running on the workstation at the time. Restarting ASDM maintains your preferences.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Tools Menu

The Tools menu provides you with troubleshooting tools on ASDM. Here you can upload new software to the ASDM, check connectivity, or issue commands at the command line.

- Command Line Interface—Provides a text-based tool for sending commands to the security appliance and viewing the results. See the Command Line Interface dialog box for more information.
- Packet Tracer—Lets you trace a packet from a specified source address and interface to a destination. You can specify the protocol and port of any type of data and see the lifespan of a packet with detailed information about actions taken on it. See the Packet Tracer dialog box for more information.
- Ping—Provides a useful tool for verifying the configuration and operation of the security appliance and surrounding communications links, as well as basic testing of other network devices. See the Ping dialog box for more information.
- Traceroute—Lets you determine the route packets will take to their destination. See the Traceroute dialog box for more information.
- File Management—Lets you view, move, copy and delete files stored in Flash memory. You can also create a directory in Flash memory. See the File Management dialog box for more information. You can also bring up the File Transfer dialog box to transfer files between various file systems, including TFTP, Flash memory, and your local PC.
- Upload ASDM Assistant Guide—Lets you upload an XML file to Flash memory that contains information used in the ASDM Assistant. These files can be downloaded from Cisco.com.
- Upgrade Software—Lets you choose a security appliance image, ASDM image, or other image file on your PC, and upload it to Flash memory. See the Upload Image from Local PC dialog box for more information.
- System Reload—Lets you restart the system and reload the saved configuration into memory. See the System Reload dialog box for more information.
- IPS/CSC Password Reset—Resets the password of an installed AIP SSM or CSC SSM to the default (cisco). See the "Resetting the AIP SSM Password" section on page 35-2 and the "Restoring the Default Password" section on page 36-12 for more information.
- ASDM Java Console—Shows the Java console.

Command Line Interface

Tools > Command Line Interface > Command Line Interface

The Command Line Interface dialog box provides a text-based tool for sending commands to the security appliance and viewing the results.

Note Commands entered via the ASDM CLI tool might function differently from commands entered through a terminal connection to the security appliance.

Command Errors

If an error occurs because you entered an incorrect command, the offending command is skipped and the remaining commands are processed anyway. A message displays in the Response box to let you know what, if any, errors were encountered as well as other pertinent information.

Note Refer to the *Cisco Security Appliance Command Reference* for a list of commands. With a few exceptions, almost all CLI commands are fully supported by ASDM.

Interactive Commands

Interactive commands are not supported in the Command Line Interface dialog box. To use these commands in ASDM, use the **noconfirm** keyword if available, as follows:

```
crypto key generate rsa modulus 1024 noconfirm
```

Avoiding Conflicts with Other Administrators

Multiple administrative users can update the running configuration of the security appliance. Before using the ASDM Command Line Interface tool to make configuration changes, check for other active administrative sessions. If more than one user is configuring the security appliance at the same time, the last changes take effect. (Click the **Monitoring** tab to view other administrative sessions that are currently active on the same security appliance.)

Viewing Configuration Changes in ASDM

If you change the configuration using the Command Line Interface tool, click the **Refresh** button to view the changes in ASDM.

Prerequisites

The commands you can enter at the Command Line Interface tool depends on your user privileges. See the Authorization Tab. Review your privilege level in the status bar at the bottom of the main ASDM window to ensure you have privileges to execute privileged-level CLI commands.

Fields

- Command—Sends commands to the security appliance.
  - Single Line—Lets you enter single commands, one at a time. The most recent commands entered are listed, or you can type a new command.
  - Multiple Line—Lets you enter multiple command lines.
  - Enable context sensitive help (?)—Shows CLI help for a command if you enter a question mark (?) after it. You do not need to press enter; the help displays as soon as you type a ?.

    Clearing this check box causes ASDM to escape the question mark character before sending it to the device, allowing you to enter the question mark as part of a text string without causing the command line help to display.
- Response—Displays the results of the commands you entered in the command box.
- Send—Sends all commands to the security appliance.
- Clear Response—Clears all text displayed in the Response box.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Packet Tracer

The packet tracer tool provides packet tracing capabilities for packet sniffing and network fault isolation.

The tool provides detailed information about the packets and how they are processed by the security appliance. In the instance that a command from the configuration did not cause the packet to drop, the packet tracer tool will provide information about the cause in an easily readable manner. For example if a packet was dropped because of an invalid header validation, a message is displayed that says, "packet dropped due to bad ip header (reason)."

In addition to capturing packets, it is possible to trace the lifespan of a packet through the security appliance to see if it is behaving as expected. The packet tracer tool lets you do the following:

- Debug all packet drops in production network.
- Verify the configuration is working as intended.
- Show all rules applicable to a packet along with the CLI lines which caused the rule addition.
- Show a time line of packet changes in a data path.
- Inject tracer packets into the data path.

Fields

- **Interface**—Specifies the source interface for the packet trace.
- **Packet type**—Specifies the protocol type for the packet trace. Available protocol types are *icmp*, *rawip*, *tcp* or *udp*.
  - **Source IP**—Specifies the source address for the packet trace.
  - **Source Port**—Specifies the source port for the packet trace.
  - **Destination IP** —Specifies the destination address for the packet trace.
  - **Destination Port**—Specifies the destination port for the packet trace.
- **Start** —Starts the packet trace.
- **Clear**—Clears all fields.
- **Show animation**—Check to display graphically the packet trace.
- **Information Display Area**—Displays detailed messages about the packet trace.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Ping

Tools > Ping > Ping

The Ping dialog box provides a useful tool for verifying the configuration and operation of the security appliance and surrounding communications links, as well as basic testing of other network devices.

A ping is the network equivalent of sonar for submarines. A ping is sent to an IP address and it returns an echo, or reply. This simple process enables network devices to discover, identify, and test each other.

The Ping tool uses ICMP described in RFC-777 and RFC-792. ICMP defines an *echo* and *echo reply* transaction between two network devices, which has become known as a ping. 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*.

Using the Ping Tool

Administrators can use the ASDM Ping tool as an interactive diagnostic aid in several ways, for example:

- Loopback testing of two interfaces—A ping may be initiated from one interface to another on the same security appliance, as an external loopback test to verify basic “up” status and operation of each interface.
- Pinging to an security appliance interface—An interface on another security appliance may be pinged by the Ping tool or another source to verify that it is up and responding.
- Pinging through an security appliance—Ping packets originating from the Ping tool may pass through an intermediate security appliance 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 security appliance interface to a network device that is suspected to be functioning improperly. If the interface is configured properly 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 security appliance interface to a network device which is known to be functioning properly and returning echo requests. If the echo is received, the proper operation of any intermediate devices and physical connectivity is confirmed.

Troubleshooting the Ping Tool

When pings fail to receive an echo, it may be the result of a configuration or operational error in a security appliance, and not always due to “NO response” from the IP address being pinged. Before using the Ping tool to ping *from*, *to* or *through* an security appliance interface, verify the following:

Basic Interface Checks

- Verify that interfaces are configured properly in Configuration > Properties > Interfaces.

- Verify that devices in the intermediate communications path, such as switches or routers, are properly delivering other types of network traffic.
- Make sure that traffic of other types from "known good" sources is being passed. Use Monitoring > Interface Graphs.

Pinging from an security appliance interface

For basic testing of an interface, a ping may be initiated from an security appliance interface to a network device which, by other means, is known to be functioning properly and returning echoes via the intermediate communications path.

- Verify receipt of the ping from the security appliance interface by the "known good" device. If it is not received, there may be a problem with the transmit hardware or configuration of the interface.
- If the security appliance interface is configured properly and it does not receive an echo from the "known good" device, there may be problems with the interface hardware receive function. If a different interface with "known good" receive capability can receive an echo after pinging the same "known good" device, the hardware receive problem of the first interface is confirmed.

Pinging to an security appliance interface

When attempting to ping *to* an security appliance interface, verify that pinging response (ICMP *echo reply)*, is enabled for that interface in the Configuration > Properties > Administration > ICMP panel. When pinging is disabled, the security appliance cannot be detected by other devices or software applications, and will not respond to the ASDM Ping tool.

Pinging through the security appliance

- First, verify that other types of network traffic from "known good" sources is being passed through through the security appliance. Use Monitoring > Interface Graphs, or an SNMP management station.
- To enable internal hosts to ping external hosts, ICMP access must be configured correctly for both the inside and outside interfaces in Configuration > Access Rules.

Fields

- IP Address—The destination IP address for the ICMP echo request packets.

Note If a host name has been assigned in the **Configuration>Hosts/Networks>Basic Information>Host Name** panel, you can use the host name in place of the IP address.

- Interface—(Optional). The security appliance interface that transmits the *echo* request packets is specified. If it is not specified, the security appliance checks the routing table to find the destination address and uses the required interface.
- Ping Output—The result of the ping. When you click **Ping**, three attempts are made to ping the IP address, and three results display the following fields:
  - Reply IP address/Device name—The IP address of the device pinged or a device name, if available. The name of the device, if assigned Hosts/Networks, may be displayed, even if **NO response** is the result.
  - Response time/timeout (ms)—When the ping is transmitted, a millisecond timer starts with a specified maximum, or timeout value. This is useful for testing the relative response times of different routes or activity levels, for example.

    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)
```

- Ping—Sends an ICMP *echo* request packet from the specified or default interface to the specified IP address and starts the response timer.
- Clear Screen—Clears the output on the screen from previous ping command attempts.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Traceroute

The Traceroute dialog box provides a useful tool to determine the route packets will take to their destination.

Traceroute Output

The traceroute tool prints the result of each probe sent. Every line of output corresponds to a TTL value in increasing order. The following are the output symbols printed by the traceroute 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 unreachablee.
!H	ICMP host unreachable.
!P	ICMP protocol unreachable.
!A	ICMP administratively prohibited.
?	Unknown ICMP error.

Fields

- Hostname or IP address—Specifies the hostname of the host to which the route is traced. If the hostname is specified, define it with **Configuration > Global Objects > IP Names**, or configure a DNS server to enable traceroute to resolve the hostname to an IP address.
- Timeout—Specifies the amount of time in seconds to wait for a response before the connection times out. The default is three seconds.
- Port—Specifies the destination port used by the UDP probe messages. The default is 33434.
- Probe—Specifies the number of probes to be sent at each TTL level. The default count is 3.

- Min & Max TTL—Specifies the minimum and maximum time to live 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 tool terminates when the traceroute packet reaches the destination or when the maximum value is reached.
- Destination Port—Specifies the destination port used by the UDP probe messages. The default is 33434.
- Specify Source Interface or IP Address—Specifies the source interface or IP address for the packet trace. 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 security appliance.
- Reverse Resolve—When checked, the output displays the names of hops encountered if name resolution is configured . If left unchecked, the output displays IP addresses.
- Use ICMP—Specifies the use of ICMP probe packets instead of UDP probe packets.
- Traceroute Output—Displays detailed messages about the traceroute.
- Traceroute—Starts the traceroute.
- Clear—Clears all fields.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

File Management

Tools > File Management > File Management

Lets you view, move, copy and delete files stored on Flash memory. You can also create a directory in Flash memory.

In multiple context mode, this tool is only available in the system.

Fields

- Folders—Displays the folders available in disk.
  - Flash Space—Shows the size of Flash and how much is available.

    Total—Shows the total size of Flash memory.

    Available—Shows how much memory is available.
- Files—Displays information about the files in the selected folder.
  - Path—Shows the selected path
  - Filename
  - Size (bytes)
  - Time Modified
  - Status

- View—Displays the selected file in your browser.
- Cut—Cuts the selected file for pasting to another directory.
- Copy—Copies the selected file for pasting to another directory.
- Paste—Pastes the copied file to the selected destination.
- Delete—Deletes the selected file from Flash.
- Rename—Lets you rename the file.
- New Directory—Creates a new directory for storing files.
- File Transfer—Opens the File Transfer dialog box.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

Upload Image from Local PC

Tools > Upgrade Software > Upload Image from Local PC

The Upload Image from Local PC dialog box lets you choose a security appliance image file, ASDM image, or other images on your PC, and upload it to Flash memory.

Fields

- Image to upload—Select which image type to upload.
- Local File Path—Enter the path to the file on your PC.
  - Browse Local—Select to browse to the file on your PC.
- Flash File System Path—Enter the path to copy the file in Flash memory.
  - Browse Local—Select to browse to the directory or file in Flash memory.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

File Transfer

Tools > File Management > File Management > File Transfer

File Transfer lets you copy files to and from your security appliance using HTTPS, TFTP, FTP or by browsing for a local image.

Fields

- Source File—Select the source file to be transferred.
  - Remote Server—Select to transfer a file from a remote server.

    Path—Enter the path to the location of the file, including the IP address of the server.

    Port/Type—Enter the port number or type (if FTP) of the remote server. Valid FTP types are:

    ap—ASCII files in passive mode.

    an—ASCII files in non-passive mode.

    ip—Binary image files in passive mode.

    in—Binary image files in non-passive mode.
  - Flash File System—Select to copy the file from Flash memory.

    Path—Enter the path to the location of the file.

    Browse Flash—Select to browse to the file location on your security appliance where the file will be copied from.
  - Local Computer—Select to copy the file from the local PC.

    Path—Enter the path to the location of the file.

    Browse Localhost—Browses the local PC for the file to be transferred.
- Destination File—Select the destination file to be transferred. Depending on the source destination, the Flash File System or the Remote Server will automatically be selected.
  - Flash File System—Transfers the file to Flash memory.

    Path—Enter the path to the location of the file.

    Browse Flash—Select to browse to the file location on your security appliance where the file will be transferred.
  - Remote Server—Transfers a file to a remote server.

    Path—Enter the path to the location of the file.

    Type—For FTP transfers, enter the type. Valid types are:

    ap—ASCII files in passive mode.

    an—ASCII files in non-passive mode.

    ip—Binary image files in passive mode.

    in—Binary image files in non-passive mode.
- Transfer File—Starts the file transfer.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

Upload ASDM Assistant Guide

Tools > Upload ASDM Assistant Guide

Upload ASDM Assistant Gude lets you upload an XML file to flash that contains useful ASDM procedural help about certain tasks. You can obtain these files from Cisco.com. Once loaded the files are available in the Search field in the File Menu.

Fields

- File to upload—The name of the XML file located on your computer, typically obtained from Cisco.com
- Flash File System Path—The path in the Flash memory where the XML file is loaded.
- Upload File—Starts the upload.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

System Reload

Tools > System Reload > System Reload

System Reload lets you restart the system and reload the saved configuration into memory. The System Reload dialog box lets you choose when the system should be reloaded, whether you should save the running configuration to Flash memory, and send a message to connected users at reload.

Fields

- Reload Scheduling—Lets you configure when the reload will take place.
  - Configuration State—Select whether to save the running configuration or not at reload.

    Save the Running Configuration at Time of Reload—Select to save the running configuration at reload.

    Reload Without Saving the Running Configuration—Select to discard configuration changes to the running configuration at reload.
- Reload Start Time—Lets you select the time of the reload.
  - Now—Select to perform an immediate reload.

  - Delay by—Lets you delay the reload by a select amount of time. Enter the time to elapse before the reload in hours and minutes or minutes.
  - Schedule at—Lets you schedule the reload to take place at a specific time and date. Enter the time of day the reload is to take place, and select the date of the scheduled reload.
- Reload Message—Enter a message to be sent to open instances of ASDM at reload.
- On Reload Failure Force Immediate Reload after—If the reload fails, the amount of time elapsed in hours and minutes or minutes before a reload is attempted again.
- Schedule Reload—Schedules the reload as configured.
- Reload Status—Displays the status of the reload.
- Cancel Reload—Cancels the scheduled reload.
- Refresh—Refreshes the Reload Status display.
- Details—Displays the details of the scheduled reload.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

Wizards Menu

The Wizards menu lets you run a wizard to configure multiple features.

- Startup Wizard—The ASDM Startup Wizard walks you, step by step, through the initial configuration of your security appliance. As you click through the configuration screens, you will be prompted to enter information about your security appliance. The Startup Wizard will apply these settings, so you should be able to start using your security appliance right away.
- VPN Wizard—The VPN Wizard is a simple way to get a VPN policy configured on your security appliance.
- High Availability and Scalability Wizard—Use this wizard to get failover configured on your security appliance.

Help Menu

The Help menu provides links to online Help as well as information about ASDM and security appliance.

- Help Topics—Opens a new browser window with help arranged by contents, screen name, and indexed in the left frame. Use these to find help for any topic, or search using the Search tab above.
- Help for Current Screen—Opens context sensitive help about the screen, panel or dialog box that is currently open. You can also click the question mark help icon for context sensitive help.
- Release Notes—Opens the most current version of the *Cisco ASDM Release Notes* on the web. The Release Notes contain the latest information about ASDM software and hardware requirements, and the latest information about changes in the software.

- Getting Started—Brings up the Getting Started help topic to help you get started using ASDM.
- Glossary—Contains definitions of terms and acronyms.
- Feature Matrix—Opens the most current version of the *Cisco ASDM Release Notes* on the web, which includes the latest licensing information.
- Feature Search—Lets you search for a function in ASDM. The Search feature looks through the titles of each panel and presents you with a list of matches, and gives you a hyperlink directly to that panel. If you need to switch quickly between two different panels you found in Search, use the Back and Forward buttons. You can also click the Search icon on the ASDM Toolbar.
- How do I?—Opens the ASDM Assistant, which lets you search downloadable content with from Cisco.com, with details about performing certain tasks.
- Legend—Provides a list of icons found in ASDM and explains what they represent.
- About Cisco *Platform*—Displays an extensive list of information about the security appliance, including software versions, hardware sets, configuration file loaded at startup, and software image loaded at startup. This information is helpful in troubleshooting.
- About Cisco ASDM 5.2—Displays information about ASDM such as the ASDM software version, hostname, privilege level, operating system, browser type, and Java version.

Toolbar

The Toolbar at the top of the ASDM window, below the menus, provides access to the home page, configuration pages, and monitoring pages. It also lets you choose between the system and security contexts in multiple context mode, and provides navigation, and other commonly-used functions.

- System/Contexts—Click the down arrow to open the context list in a left-hand pane, and the up arrow to restore the context drop-down list. When expanded, click the left arrow to collapse the pane all the way left, and the right arrow to restore the pane. To manage the system, select System from the list. To manage a context, select the context from the list.
- Home—Displays the Home page, which lets you view at a glance important information about your security appliance such as the status of your interfaces, the version you are running, licensing information, and performance. See Home Page for more information. In multiple mode, the system does not have a Home page.
- Configuration—Configures the security appliance. Choose a feature button in the left-hand pane to configure that feature.
- Monitoring—Monitors the security appliance. Choose a feature button in the left-hand pane to monitor that feature.
- Back—Takes you back to the last panel of ASDM you visited.
- Forward—Takes you forward to the last panel of ASDM you visited.
- Search—Lets you search for a function in ASDM. The Search feature looks through the titles of each panel and presents you with a list of matches, and gives you a hyperlink directly to that panel. If you need to switch quickly between two different panels you found in Search, use Back and Forward.
- Refresh—Refreshes ASDM with the current running configuration by selecting. This button does not refresh the graphs in any of the monitoring graphs.
- Save—Saves the running configuration to the startup configuration. If you have a context that is not write accessible, for example on HTTP, then this button does not save the running configuration.
- Help—Shows context-sensitive help for the screen that is currently open.

Status Bar

The status bar appears at the bottom of the ASDM window. The areas below appear from left to right on the status bar.

- Status—Shows the status of the configuration, such as "Device configuration loaded successfully."
- User Name—Shows the username of the ASDM user. If you logged in without a username, the username is "admin."
- User Privilege—Shows the privilege of the ASDM user.
- Commands Ignored by ASDM—When you click the icon, ASDM shows a list of commands from your configuration that ASDM did not process. They will not be removed from the configuration. See Show Commands Ignored by ASDM on Device for more information.
- Status of Connection to Device—Shows the ASDM connection status to the security appliance. See Connection to Device for more information.
- Save to Flash Needed—Shows that you made configuration changes in ASDM, but that you have not yet saved the running configuration to the startup configuration.
- Refresh Needed—Shows that you need to refresh the configuration from the security appliance to ASDM because the configuration changed on the security appliance. For example, you made a change to the configuration at the CLI.
- SSL Secure—Shows that the connection to ASDM is secure because it uses SSL.
- Time—Shows the time that is set on the switch that contains the security appliance.

Connection to Device

Status Bar > Status of Connection to Device icon > Connection to Device

ASDM maintains a constant connection to the security appliance to maintain up-to-date monitoring and home page data. This dialog box shows the status of this connection. When you make a configuration change, ASDM opens a second connection for the duration of the configuration, and then closes it. That connection is not represented by this dialog box.

Buttons That Appear on Many Panels

These buttons appear on many ASDM panels:

- Apply—Sends changes made in ASDM to the security appliance and applies them to the running configuration. Click **Save** to write a copy of the running configuration to Flash memory. Use the File menu to write a copy of the running configuration to Flash memory, a TFTP server, or a failover standby unit.
- **Reset**—Discards changes and reverts to the information displayed before changes were made or the last time you clicked **Refresh** or **Apply**. After Reset, use Refresh to make sure that information from the current running configuration is displayed.
- **Cancel**—Discards changes and returns to the previous panel.
- **Help**—Displays help for the selected panel.

About the Help Window

This section contains the following topics:

Header Buttons

Use the header buttons to navigate through the help to find the topic you are looking for.

- About ASDM—Displays information about ASDM.
- Search—Lets you search the help topics.
- Using Help—Describes the best way to get the most out of online help.
- Glossary—Lists a glossary of terms found in ASDM and networking.

Left-Pane Tabs—Help navigate the online help.

- Contents—Displays a table of contents.
- Screens—Lists help files by screen name.
- Index—Provides an index of help topics found in ASDM online help

Right-Pane Help Content—Displays the help for the selected topic.

Notes

When help is invoked in applet mode and if there is any help page already open, the new help page will appear in the same browser window. If there is no help page already open, then the help page will appear in a new browser window.

When help is invoked in application mode and if Netscape is the default browser, each time help is invoked the help page will appear in a new browser window. If IE is the default browser, based on the user setting, the help page may appear either in the last visited browser window or in a new browser window. This behavior of IE can be controlled by using the option Tools > Internet Options > Advanced > Reuse window for launching shortcuts.

Home Page

The ASDM home pane lets you view, at a glance, important information about your security appliance. If you have an SSM installed in your security appliance, an additional tab appears on the home page. The additional tab displays status information about the software on the SSM.

For more information about configuring these areas, see the following:

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Home

Home

The ASDM home pane lets you view, at a glance, important information about your security appliance, such as the status of your interfaces, the version you are running, licensing information, and performance.

Many of the details available on the ASDM home page are available elsewhere in ASDM, but this is a useful and quick way to see how your security appliance is running. Status information on the Home pane is updated every ten seconds.

Fields

- Device Information—Includes two tabs to show device information.
  - General—Shows the following information:

    Host Name—*Display only.* Shows the security appliance hostname. See Device to set the hostname.

    Platform Version—*Display only.* Shows the security appliance software version.

    Device Uptime—*Display only.* Shows how long the security appliance has been running.

    ASDM Version—*Display only.* Shows the ASDM version.

    Device Type—*Display only.* Shows the security appliance model.

    Firewall Mode—*Display only.* Shows the firewall mode, either Routed or Transparent. See Firewall Mode Overview for more information.

    Context Mode—*Display only.* Shows the context mode, either Single or Multiple. See Security Context Overview for more information.

    Total Flash—*Display only.* Shows the total amount of Flash memory (the internal Flash memory plus the external Flash memory card, if available) in MB.

    Total Memory—*Display only.* Shows the total RAM.
  - License—*Display only.* Shows the level of support for licensed features on the security appliance.
- VPN Status—Routed, single mode only. Shows the following information:
  - IKE Tunnels—*Display only.* Shows the number of connected IKE tunnels.
  - IPSec Tunnels—*Display only.* Shows the number of connected IPSec tunnels.
- System Resources Status—Shows the following CPU and memory usage statistics:
  - CPU—*Display only.* Shows the current percentage of CPU being utilized.
  - CPU Usage (percent)—*Display only.* Shows the CPU usage for the last five minutes.
  - Memory—*Display only.* Shows the current amount of memory being used in MB.

  - Memory Usage (MB)—*Display only.* Shows the memory usage for the last five minutes in MB.
- Interface Status—Shows the status of each interface. If you select an interface row, the input and output Kbps shows under the table.
  - Interface—*Display only.* Shows the interface name.
  - IP Address/Mask—*Display only.* Routed mode only. Shows the IP address and subnet mask of the interface.
  - Line—*Display only.* Shows the administrative status of the interface. A red icon is displayed if the line is down, and a green icon is displayed if the line is up.
  - Link—*Display only.* Shows the link status of the interface. A red icon is displayed if the link is down, and a green icon is displayed if the link is up.
  - Current Kbps—*Display only.* Shows the current number of kilobits per second that cross the interface.
- Traffic Status—Shows graphs for connections per second for all interfaces and for the traffic throughput of the lowest security interface.
  - Connections per Second Usage—*Display only.* Shows the UDP and TCP connections per second over the last 5 minutes. This graph also shows the current number of connections by type, UDP, TCP, and Total.
  - *Name* Interface Traffic Usage (Kbps)—*Display only.* Shows the traffic throughput for the lowest security interface. If you have multiple interfaces at the same level, then ASDM shows the first interface alphabetically. This graph also shows the current throughput by type, Input Kbps and Output Kbps.
- Latest ASDM Syslog Messages—Shows the latest system messages generated by the security appliance.
  - Stop Message Display—Stops logging to ASDM.
  - Resume Message Display—Resumes logging to ASDM.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Home > Content Security Tab

Home > Content Security Tab

The Content Security tab lets you view important information about the Content Security and Control (CSC) SSM. This panel appears only if a CSC SSM is installed in the security appliance.

For an introduction to CSC SSM, see About the CSC SSM.

Note If you have not completed the Setup Wizard in Configuration > Trend Micro Content Security > CSC Setup, you cannot access the panels under Home > Content Security. Instead, a dialog box appears and lets you access the Setup Wizard directly from Home > Content Security.

Fields

- Device Information—Shows the following information:

  Model—Shows the type of SSM installed in your security appliance.

  Mgmt IP—Shows the IP address of the management interface for the CSC SSM.

  Version—Shows the CSC SSM software version.

  Last Update—Shows the date of the last software update obtained from Trend Micro.

  Daily Node #—Shows the number of network devices for which the CSC SSM provided services in the preceding 24 hours. ASDM updates this field at midnight.

  Base License—Shows license status for basic features of the CSC SSM, such as anti-virus, anti-spyware, and FTP file blocking features. The date that the license is due to expire appears. If the license has expired, the date of expiry appears. If no license is configured, the field shows Not Available.

  Plus License—Shows license status for advanced features of the CSC SSM, such as anti-spam, anti-phishing, email content filtering, and URL blocking and filtering features. The date that the license is due to expire appears. If the license has expired, the date of expiry appears. If no license is configured, the field shows Not Available.

  Licensed Nodes—Shows the maximum number of network devices for which your CSC SSM is licensed to provide services.

- System Resources Status—Shows the following CPU and memory usage statistics for the CSC SSM:
  - CPU—Shows the current percentage of CPU being utilized.
  - CSC SSM CPU Usage (percent)—Shows the CPU usage for the last five minutes.
  - Memory—Shows the current amount of memory being used in MB.
  - CSC SSM Memory Usage (MB)—Shows the memory usage for the last five minutes in MB.
- Threat Summary—Shows aggregate data about threats detected by the CSC SSM.
  - Threat Type—Lists four threat types: Virus, Spyware, URL Filtered, and URL Blocked.
  - Today—Shows the number of threats detected for each threat type within the past 24 hours.
  - Last 7 Days—Shows the number of threats detected for each threat type within the past 7 days.
  - Last 30 Days—Shows the number of threats detected for each threat type within the past 30 days.
- Email Scan—Shows graphs for emails scanned and email virus and spyware detected.
  - Email Scanned Count—Shows the number of emails scanned, as separate graphs by email protocol (SMTP or POP3) and as a combined graph for both supported email protocols. The graphs display data in ten-second intervals.
  - Email Virus and Spyware—Shows the number of viruses and emails detected in email scans, as separate graphs by threat type (virus or spyware). The graphs display data in ten-second intervals.

- Latest CSC Security Events—Shows, in real time, security event messages received from the CSC SSM.
  - Time—Displays the time an event occurred.
  - Source—Displays the IP address or hostname from which the threat came.
  - Threat/Filter—Displays the type of threat or, in the case of a URL filter event, the filter that triggered the event.
  - Subject/File/URL—Displays the subject of emails containing a threat, the names of FTP file containing a threat, or URLs blocked or filtered.
  - Receiver/Host—Displays the recipient of emails containing a threat or the IP address or hostname of a node threatened.
  - Sender—Displays the sender of emails containing a threat.
  - Content Action—Displays the action taken upon the content of the message or file, such as delivering the content unaltered, deleting attachments, or cleaning attachments before delivering them.
  - Msg Action—Displays the action taken upon the message, such as delivering the message unchanged, delivering the message after deleting attachments, or not delivering the message.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Before You Start

This section contains the following topics:

Factory Default Configurations

The factory default configuration is the configuration applied by Cisco to new security appliances. The factory default configuration is supported on all models except for the PIX 525 and PIX 535 security appliances.

For the PIX 515/515E and the ASA 5510 and higher security appliances, the factory default configuration configures an interface for management so you can connect to it using ASDM, with which you can then complete your configuration.

For the ASA 5505 adaptive security appliance, the factory default configuration configures interfaces and NAT so that the security appliance is ready to use in your network immediately.

The factory default configuration is available only for routed firewall mode and single context mode. See Configuring Security Contexts for more information about multiple context mode. See the Firewall Mode Overview for more information about routed and transparent firewall mode.

This section includes the following topics:

Restoring the Factory Default Configuration

To restore the factory default configuration, perform the following steps:

Step 1 Choose **File > Reset Device to the Factory Default Configuration**.

Step 2 To change the default IP address to an IP address of your choosing, check **Use this address** for the <default interface> which will be named as <name> check box.

Step 3 Enter the new IP address in the Management IP Address field.

Step 4 Enter the new subnet mask in the Management Mask field.

Step 5 Click **OK**.

If you specify the *ip_address*, then you set the inside or management interface IP address, depending on your model, instead of using the default IP address of 198.168.1.1. The **http** command uses the subnet you specify. Similarly, the **dhcpd address** command range consists of addresses within the subnet that you specify.

After you restore the factory default configuration, save it to internal Flash memory using the **write memory** command. The **write memory** command saves the running configuration to the default location for the startup configuration, even if you previously configured the **boot config** command to set a different location; when the configuration was cleared, this path was also cleared. See the

Note This command also clears the **boot system** command, if present, along with the rest of the configuration. The **boot system** command lets you boot from a specific image, including an image on the external Flash memory card. The next time you reload the security appliance after restoring the factory configuration, it boots from the first image in internal Flash memory; if you do not have an image in internal Flash memory, the security appliance does not boot.

To configure additional settings that are useful for a full configuration, see the **setup** command.

ASA 5505 Default Configuration

The default factory configuration for the ASA 5505 adaptive security appliance configures the following:

- An inside VLAN 1 interface that includes the Ethernet 0/1 through 0/7 switch ports. If you did not set the IP address in the **configure factory-default** command, then the VLAN 1 IP address and mask are 192.168.1.1 and 255.255.255.0.
- An outside VLAN 2 interface that includes the Ethernet 0/0 switch port. VLAN 2 derives its IP address using DHCP.
- The default route is also derived from DHCP.
- All inside IP addresses are translated when accessing the outside using interface PAT.
- By default, inside users can access the outside with an access list, and outside users are prevented from accessing the inside.
- The DHCP server is enabled on the security appliance, so a PC connecting to the VLAN 1 interface receives an address between 192.168.1.2 and 192.168.1.254.

- The HTTP server is enabled for ASDM and is accessible to users on the 192.168.1.0 network.

The configuration consists of the following commands:

```
interface Ethernet 0/0
   switchport access vlan 2
   no shutdown
interface Ethernet 0/1
   switchport access vlan 1
   no shutdown
interface Ethernet 0/2
   switchport access vlan 1
   no shutdown
interface Ethernet 0/3
   switchport access vlan 1
   no shutdown
interface Ethernet 0/4
   switchport access vlan 1
   no shutdown
interface Ethernet 0/5
   switchport access vlan 1
   no shutdown
interface Ethernet 0/6
   switchport access vlan 1
   no shutdown
interface Ethernet 0/7
   switchport access vlan 1
   no shutdown
interface vlan2
   nameif outside
   no shutdown
   ip address dhcp setroute
interface vlan1
   nameif inside
   ip address 192.168.1.1 255.255.255.0
   security-level 100
   no shutdown
global (outside) 1 interface
nat (inside) 1 0 0
http server enable
http 192.168.1.0 255.255.255.0 inside
dhcpd address 192.168.1.2-192.168.1.254 inside
dhcpd auto_config outside
dhcpd enable inside
logging asdm informational
```

ASA 5510 and Higher Default Configuration

The default factory configuration for the ASA 5510 and higher adaptive security appliance configures the following:

- The management Management 0/0 interface. If you did not set the IP address in the **configure factory-default** command, then the IP address and mask are 192.168.1.1 and 255.255.255.0.
- The DHCP server is enabled on the security appliance, so a PC connecting to the interface receives an address between 192.168.1.2 and 192.168.1.254.
- The HTTP server is enabled for ASDM and is accessible to users on the 192.168.1.0 network.

The configuration consists of the following commands:

```
interface management 0/0
```

```
   ip address 192.168.1.1 255.255.255.0
   nameif management
   security-level 100
   no shutdown
asdm logging informational 100
asdm history enable
http server enable
http 192.168.1.0 255.255.255.0 management
dhcpd address 192.168.1.2-192.168.1.254 management
dhcpd lease 3600
dhcpd ping_timeout 750
dhcpd enable management
```

PIX 515/515E Default Configuration

The default factory configuration for the PIX 515/515E security appliance configures the following:

- The inside Ethernet1 interface. If you did not set the IP address in the **configure factory-default** command, then the IP address and mask are 192.168.1.1 and 255.255.255.0.
- The DHCP server is enabled on the security appliance, so a PC connecting to the interface receives an address between 192.168.1.2 and 192.168.1.254.
- The HTTP server is enabled for ASDM and is accessible to users on the 192.168.1.0 network.

The configuration consists of the following commands:

```
interface ethernet 1
   ip address 192.168.1.1 255.255.255.0
   nameif management
   security-level 100
   no shutdown
asdm logging informational 100
asdm history enable
http server enable
http 192.168.1.0 255.255.255.0 management
dhcpd address 192.168.1.2-192.168.1.254 management
dhcpd lease 3600
dhcpd ping_timeout 750
dhcpd enable management
```

Configuring the Security Appliance for ASDM Access

If you want to use ASDM to configure the security appliance instead of the command-line interface, you can connect to the default management address of 192.168.1.1 (if your security appliance includes a factory default configuration. See the "Factory Default Configurations" section on page 2-1.). On the ASA 5510 and higher adaptive security appliances, the interface to which you connect with ASDM is Management 0/0. For the ASA 5505 adaptive security appliance, the switch port to which you connect with ASDM is any port, except for Ethernet 0/0. For the PIX 515/515E security appliance, the interface to which you connect with ASDM is Ethernet 1.

If you do not have a factory default configuration, see the *Cisco Security Appliance Command Line Configuration Guide* to access the command-line interface. You can then configure the minimum parameters to access ASDM by entering the **setup** command.

Setting Transparent or Routed Firewall Mode at the CLI

You can set the security appliance to run in routed firewall mode (the default) or transparent firewall mode. For more information about the firewall mode, see Firewall Mode Overview.

For multiple context mode, you can use only one firewall mode for all contexts. You must set the mode in the system execution space.

When you change modes, the security appliance clears the configuration because many commands are not supported for both modes. If you already have a populated configuration, be sure to back up your configuration before changing the mode; you can use this backup for reference when creating your new configuration. For multiple context mode, the system configuration is erased. This action removes any contexts from running. If you then re-add a context that has an existing configuration that was created for the wrong mode, the context configuration will not work correctly. Be sure to recreate your context configurations for the correct mode before you re-add them, or add new contexts with new paths for the new configurations.

If you download a text configuration to the security appliance that changes the mode with the **firewall transparent** command, be sure to put the command at the top of the configuration; the security appliance changes the mode as soon as it reads the command and then continues reading the configuration you downloaded. If the command is later in the configuration, the security appliance clears all the preceding lines in the configuration.

To set the firewall mode, perform the following steps. In multiple context mode, perform these steps in the system execution space.

Step 1 In single context mode or from the system configuration in multiple mode, you can copy the startup configuration or running configuration to an external server or to the local Flash memory using one of the following commands. You can use this backup configuration for reference when creating your new configuration.

- To copy to a TFTP server, enter the following command:

  ```
  hostname# copy {startup-config | running-config} tftp://server[/path]/filename
  ```

- To copy to a FTP server, enter the following command:

  ```
  hostname# copy {startup-config | running-config}
  ftp://[user[:password]@]server[/path]/filename
  ```

- To copy to local Flash memory, enter the following command:

  ```
  hostname# copy {startup-config | running-config} {flash:/ | disk0:/ |
  disk1:/}[path/]filename
  ```

  Be sure the destination directory exists. If it does not exist, first create the directory using the **mkdir** command.

Step 2 To change the mode, enter one of the following commands:

- To set the mode to transparent, enter the following command:

  ```
  hostname(config)# firewall transparent
  ```

  This command also appears in each context configuration for informational purposes only; you cannot enter this command in a context.

- To set the mode to routed, enter the following command:

  ```
  hostname(config)# no firewall transparent
  ```

Downloading the ASDM Launcher

The ASDM Launcher is for Windows only. The ASDM Launcher is an improvement over running ASDM as a Java Applet. The ASDM Launcher avoids double authentication and certificate dialog boxes, launches faster, and caches previously-entered IP addresses and usernames.

To download the ASDM launcher, perform the following steps:

Step 1 From a supported web browser on the security appliance network, enter the following URL:

https://*interface_ip_address*

In transparent firewall mode, enter the management IP address.

Note Be sure to enter **https**, not **http**.

Step 2 Click **OK** or **Yes** to all prompts, including the name and password prompt. By default, leave the name and password blank.

A page displays with the following buttons:

- **Download ASDM Launcher and Start ASDM**
- **Run ASDM as a Java Applet**

Step 3 Click **Download ASDM Launcher and Start ASDM**.

The installer downloads to your PC.

Step 4 Run the installer to install the ASDM Launcher.

Starting ASDM

This section describes how to start ASDM according to one of the following methods:

- Starting ASDM from the ASDM Launcher, page 2-6
- Using ASDM in Demo Mode, page 2-7
- Starting ASDM from a Web Browser, page 2-8

Starting ASDM from the ASDM Launcher

The ASDM Launcher is for Windows only.

To start ASDM from the ASDM Launcher, perform the following steps:

Step 1 Double-click the Cisco ASDM Launcher shortcut on your desktop, or start it from the **Start** menu.

Step 2 Enter the security appliance IP address or hostname, your username, and your password, and then click **OK**.

If there is a new version of ASDM on the security appliance, the ASDM Launcher automatically downloads it before starting ASDM.

Using ASDM in Demo Mode

ASDM Demo Mode is available as a separately installed application running under Windows. It makes use of the ASDM Launcher and pre-packaged configuration files to let you run ASDM without having a live device available. ASDM Demo Mode lets you:

- Perform configuration and select monitoring tasks via ASDM as though you were interacting with a real device.
- Demonstrate ASDM or security appliance features using the ASDM interface.
- Perform configuration and monitoring tasks with the Content Security and Control SSM (CSC SSM).

ASDM Demo Mode provides simulated monitoring data, including real-time system log messages. The data shown is randomly generated, but the experience is identical to what you would see when connecting to a real device.

ASDM Demo Mode has the following limitations:

- Changes made to the configuration will appear in the GUI but are not applied to the configuration file. That is, when you click the Refresh button, it will revert back to the original configuration. The changes are never saved to the configuration file.
- File/Disk operations are not supported.
- Monitoring and logging data are simulated. Historical monitoring data is not available.
- You can only log in as an admin user; you cannot login as a monitor-only or read-only user.
- Demo Mode does not support the following features:
  - File menu:

    Save Running Configuration to Flash

    Save Running Configuration to TFTP Server

    Save Running Configuration to Standby Unit

    Save Internal Log Buffer to Flash

    Clear Internal Log Buffer
  - Tools menu:

    Command Line Interface

    Ping

    File Management

    Update Image

    File Transfer

    Upload image from Local PC

    System Reload

  - Toolbar/Status bar > Save
  - Configuration > Interface > Edit Interface > Renew DHCP Lease
  - Failover—Configuring a standby device
- These operations cause a reread of the configuration and therefore will revert it back to the original configuration.
  - Switching contexts
  - Making changes in the Interface panel
  - NAT panel changes
  - Clock panel changes

To run ASDM in Demo Mode, perform the following steps:

Step 1 If you have not yet installed the Demo Mode application, perform the following steps:

a. Download the ASDM Demo Mode installer from http://www.cisco.com/cgi-bin/tablebuild.pl/cat6000-fwsm.

   The filename is asdm-*version*-demo.msi.

b. Double-click the installer to install the software.

Step 2 Double-click the Cisco ASDM Launcher shortcut on your desktop, or start it from the **Start** menu.

Step 3 Click the **Run in Demo Mode** check box.

Step 4 To set the platform, context and firewall modes, and ASDM Version, click the **Demo** button and make your selections from the Demo Mode area.

Step 5 If you want to use new ASDM images as they come out, you can either download the latest installer, or you can download the normal ASDM images and install them for Demo Mode:

a. Download the image from http://www.cisco.com/cgi-bin/tablebuild.pl/cat6000-fwsm.

   The filename is asdm-*version*.bin

b. In the Demo Mode area, click **Install ASDM Image**.

   A file browser appears. Find the ASDM image file in the browser.

Step 6 Click **OK** to launch ASDM Demo Mode.

You see a Demo Mode label in the title bar of the window.

Starting ASDM from a Web Browser

To start ASDM from a web browser, perform the following steps:

Step 1 From a supported web browser on the security appliance network, enter the following URL:

https://*interface_ip_address*

In transparent firewall mode, enter the management IP address.

Note Be sure to enter **https**, not **http**.

Step 2 Click **OK** or **Yes** to all browser prompts, including the name and password prompt. By default, leave the name and password blank.

A page displays with the following buttons:

- **Download ASDM Launcher and Start ASDM**
- **Run ASDM as a Java Applet**

Step 3 Click **Run ASDM as a Java Applet**.

Step 4 Click **OK** or **Yes** to all Java prompts, including the name and password prompt. By default, leave the name and password blank.

History Metrics

Configuration > Properties > History Metrics

The History Metrics pane lets you configure the security appliance to keep a history of various statistics, which can be displayed by ASDM on any Graph/Table. If you do not enable history metrics, you can only monitor statistics in real time. Enabling history metrics lets you view statistics graphs from the last 10 minutes, 60 minutes, 12 hours, and 5 days.

Fields

- ASDM History Metrics—Enables history metrics. Unchecking this check box clears and disables the history metrics.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Configuration Overview

To configure and monitor the security appliance, perform the following steps:

Step 1 Use the Startup Wizard for initial configuration by clicking **Wizards > Startup Wizard**.

Step 2 To configure VPN connections, use the VPN Wizard by clicking **Wizards > VPN Wizard** and completing each screen that appears.

Step 3 Configure advanced features by clicking the **Configuration** button on the toolbar and then clicking a feature button. Features include:

- Configuring Interfaces—Configures basic interface parameters including the IP address, name, security level, and for transparent mode, the bridge group.
- Security Policy—Includes access rules, AAA rules, filter rules, and service policy rules.

  - Access Rules—Permits or denies IP traffic through the security appliance. For transparent firewall mode, you can also apply an EtherType access list to allow non-IP traffic.
  - Ethertype Rules (Transparent Mode Only)—Permits or denies non-IP traffic through the security appliance.
  - AAA Rules—Requires authentication and/or authorization for certain types of traffic, for example, for HTTP. The security appliance also sends accounting information to a RADIUS or TACACS+ server.
  - Filter Rules—Prevents outbound access to specific websites or FTP servers. The security appliance works with a separate server running either Websense Enterprise or Sentian by N2H2. See Configuration > Properties > URL Filtering to configure the URL filtering server, which must be configured before you add a rule.
  - Service Policy Rules—Applies application inspection, connection limits, and TCP normalization. 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 security appliance to do a deep packet inspection. You can also limit TCP and UDP connections and embryonic connections. Limiting the number of connections and embryonic connections protects you from a DoS attack. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. TCP normalization drops packets that do not appear normal.

- NAT—Translates addresses used on a protected network to addresses used on the public Internet. This lets you use private addresses, which are not routable on the Internet, on your inside networks.
- VPN—Configures VPN connections.
  - VPN Wizard—Runs the VPN wizard.
  - E-Mail Proxy—Configures e-mail proxies. E-mail proxies extend remote e-mail capability to WebVPN users.
  - General—Sets general VPN configuration parameters.
  - IKE—IKE, also called ISAKMP, is the negotiation protocol that lets two hosts agree on how to build an IPSec security association.
  - IP Address Management—Sets the IP addresses of clients after they connect through the VPN tunnel.
  - IPSec—Configures the IPSec protocol for VPN tunnels.
  - Load Balancing—Configures load balancing for VPN connections.
  - WebVPN—Configures WebVPN. WebVPN lets users establish a secure, remote-access VPN tunnel to the security appliance using a web browser.
- CSD Manager—Configures the CSC SSM (available for the ASA 5500 series adaptive security appliance).
- Configuring IPS—Configures the AIP SSM (available for the ASA 5500 series adaptive security appliance).
- Configuring Dynamic And Static Routing—(Single mode only) Configures OSPF, RIP, static, and asymmetric routing.
- Global Objects—Provides a single location where you can configure, view, and modify the reusable components that you need to implement your policy on the security appliance. These reusable components, or global objects, include the following:
  - Hosts/Networks
  - Inspect Maps

  - TCP Maps
  - Time Ranges

Step 4 Monitor the security appliance by clicking the **Monitoring** button on the toolbar and then clicking the feature button. Features include:

- Monitoring Interfaces—Monitors the ARP table, DHCP, dynamic access list, and interface statistics.
- Monitoring Routing—Monitors routes, OSPF LSAs, and OSPF neighbors.
- Monitoring Properties—Monitors management sessions, AAA servers, failover, CRLs, the DNS cache, and system statistics.
- Monitoring System Log Messages—Monitors system log messages.
- Monitoring Failover—(For the system in multiple mode) Monitors failover in the system.

Using the Startup Wizard

Startup Wizard

From Cisco ASDM 5.2 Start Screen

Configuration > Properties > Startup Wizard

The ASDM Startup Wizard walks you through, step by step, the initial configuration of your security appliance. As you click through the configuration screens, you will be prompted to enter information about your security appliance. The Startup Wizard will apply these settings, so you should be able to start using your security appliance right away.

The Startup Wizard defines the following in your configuration:

- A hostname for your security appliance.
- A domain name for your security appliance.
- An enable password that is used to restrict administrative access to the security appliance through ASDM or the command-line interface.
- The IP address information of the outside interface on the security appliance.
- The other interfaces of your security appliance, such as the inside or DMZ interfaces, can be configured from the Startup Wizard.
- NAT or PAT rules for your security appliance.
- DHCP settings for the inside interface, such as for use with a DHCP server.

More information about each setting is available by clicking the Help button on the corresponding configuration screen.

Before you begin using the Startup Wizard, make sure you have the following information available:

- A unique hostname to identify the security appliance on your network.
- The IP addresses of your outside, inside, and other interfaces.
- The IP addresses to use for NAT or PAT configuration.
- The IP address range for the DHCP server.

Remember:

- You can access the Startup Wizard from the Cisco ASDM 5.2 Start page by selecting the ‘Run Startup Wizard as a Java Applet’ button.
- You can access the Startup Wizard at any time using the Wizards menu in ASDM.
- The Help button is an icon with a question mark.

- On subsequent Startup Wizard pages, you can click **Finish** to complete the wizard at any time. This sends changes made in the Startup Wizard to the security appliance.

Important Notes

- The security appliance can run in two modes:
  - Routed—In routed mode, the security appliance acts as a router between connected networks. Each interface requires an IP address on a different subnet. The security appliance performs NAT between connected networks. In single context mode, the routed firewall supports OSPF and RIP (in passive mode). Multiple context mode supports static routes only. Routed mode supports up to 256 interfaces per context or in single mode, with a maximum of 1000 interfaces divided between all contexts. Each interface is on a different subnet. You can share interfaces between contexts.

    In routed mode, some types of traffic cannot pass through the security appliance even if you allow it in an ACL. The transparent firewall, however, can allow any traffic through using either an extended ACL (for IP traffic) or an EtherType ACL.

    Note We recommend using the advanced routing capabilities of the upstream and downstream routers, such as the MSFC, instead of relying on the security appliance for extensive routing needs.

  - Transparent—In transparent mode, the security appliance is not seen as a router hop to connected devices, but acts like a "bump in the wire," or a "stealth firewall." The security appliance connects the same network on its inside and outside ports, but uses different VLANs on the inside and outside. No dynamic routing protocols or NAT are required. Transparent mode only supports two interfaces, an inside interface and an outside interface. Transparent mode helps simplify the configuration and reduces its visibility to attackers. You can also use a transparent firewall for traffic that would otherwise be blocked in routed mode. For example, a transparent firewall can allow multicast streams.

    Even though transparent mode acts as a bridge, Layer 3 traffic, such as IP traffic, cannot pass through the security appliance unless you explicitly permit it with an extended ACL. The only traffic allowed through the transparent firewall without an ACL is ARP traffic. ARP traffic can be controlled by ARP inspection.

    Note The transparent mode security appliance does not pass CDP packets.

- The context mode (single or multiple) is not stored in the configuration file, even though it does endure reboots. If you need to copy your configuration to another device, set the mode on the new device to match using the mode command.

  When you convert from single mode to multiple mode, the security appliance converts the running configuration into two files: a new startup configuration that comprises the system configuration, and admin.cfg that comprises the admin context (in the root directory of the internal Flash memory).

- With a full license, the security appliance supports up to five interfaces with a maximum of three interfaces named 'interface.' In restricted mode, the security appliance supports up to three interfaces, and in transparent mode, the security appliance supports up to two interfaces.

  Once the maximum number of interfaces has been created, or the maximum number of interfaces has already been named, you may not be able to create a new VLAN, and may have to select from an existing VLAN.

- The security appliance can be used as an Easy VPN remote device. However, if the security appliance is configured to function as an Easy VPN remote device, it cannot establish other types of tunnels. For example, the security appliance cannot function simultaneously as both an Easy VPN remote device and as one end of a standard peer-to-peer VPN deployment.

  There are two modes of operation when using the security appliance as an Easy VPN remote device:

  – Client Mode

  – Network Extension Mode

  When configured in Easy VPN Client Mode, the security appliance does not expose the IP addresses of clients on its inside network. Instead, it uses NAT (Network Address Translation) to translate the IP addresses on the private network to a single, assigned IP address. When thesecurity appliance is configured in Client Mode, you cannot ping or access any device from outside the private network.

  When configured in Easy VPN Network Extension Mode, the security appliance does not protect the IP addresses of local hosts by substituting a assigned IP address. Therefore, hosts on the other side of the VPN connection can communicate directly with hosts on the local network.

Fields

Launch Startup Wizard—Launches the Startup Wizard.

Note The Launch Startup Wizard button does not appear if you click Wizards >Startup Wizard on the toolbar.

With the exception of this screen, all screens in the Startup Wizard display the following buttons:

- Back—Returns you to the previous screen (the button is dim in this screen).
- Next—Advances you to the next screen.
- Finish—Submits your configuration to the security appliance based upon choices made in this screen (the button is dim in this screen).
- Cancel—Discards any changes without applying them. The Wizard prompts you with the Exit Wizard dialog box when Cancel is clicked. Clicking Exit closes the Wizard, and clicking Cancel again returns you to the Wizard screen. Remember at any time in the Wizard you can click Back to return to the previous screen.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Starting Point

Configuration >Properties >Startup Wizard > Starting Point

Benefits

The Starting Point screen lets you continue with your existing configuration or reset the configuration to the factory default values. If you check the box 'Reset configuration to existing defaults,' you revert back to the IP address and subnet mask of the default inside interface. If you continue with your existing configuration, you automatically retain your IP address and subnet mask.

Fields

The Starting Point screen displays the Next, Cancel, and Help buttons, in addition to the following:

- Modify existing configuration—Click to start the wizard with the existing configuration.
- Reset configuration to factory defaults—Click to start the wizard at the factory default values for the inside interface.
  - Configure the IP address of the management interface—Check this box to configure the IP address and subnet mask of the management interface.

    IP Address—Lets you enter the IP address of the management interface to configure.

    Subnet Mask—Lets you enter the subnet mask of the management interface to configure.

Note If you reset the configuration to the factory defaults, you cannot undo the change by cancelling the wizard.

Note The Back and Finish buttons are disabled on this screen.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
—	—	•	—	—

For More Information

This feature is available in the main ASDM application screen:

File > Reset Device to the Factory Default Configuration

Basic Configuration

Startup Wizard >Basic Configuration

Benefits

The Basic Configuration screen lets you configure the hostname of your security appliance and the enable password, as well as a domain name for the security appliance.

The domain name should be less than 64 characters (maximum 63 characters) alphanumeric and mixed case.

The enable password is used to administer ASDM or to administer the security appliance from the Command Line Interface. The password is case-sensitive and can be up to 16 alphanumeric characters. If you want to change the current password, check **Change privileged mode (enabled) password**, enter the old password, then enter the new password, and confirm the new password in the fields provided.

Note If you leave the password field blank, a Password Confirmation screen displays and notifies you that this is a high security risk.

Fields

The Basic Configuration screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Host Name—Lets you enter a hostname for the security appliance. The hostname can be up to 63 alphanumeric characters and mixed case. Note: This field will list either ASA or PIX before Host Name, depending on the platform you are using.
- Domain Name—Specifies the IPSec domain name the of the security appliance. This can be used later for certificates. There is a 64-character limit on the domain name (maximum 63 characters), and it must be alphanumeric with no special characters or spaces.
- Privileged Mode (Enable) Password area—Lets you restrict administrative access to the security appliance through ASDM or the Command Line Interface.
  - Change privileged mode (enable) Password—Check this box to change the current privileged mode (enable) password.
  - Old Password—Lets you enter the old enable password, if one exists.
  - New Password—Lets you enter the new enable password. The password is case-sensitive and can be up to 16 alphanumeric characters.
  - Confirm New Password—Lets you reenter the new enable password.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > Properties > Device Administration > Device

Configuration > Properties > Device Administration > Password

Outside Interface Configuration

Startup Wizard > Outside Interface Configuration

Benefits

The Outside Interface Configuration screen lets you configure your outside interface by specifying an IP address, or obtaining one from a PPPoE or a DHCP server.

Fields

The Outside Interface Configuration screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Interface Properties area
  - Interface—Lets you add a new interface, or select an interface from the drop-down list.
  - Interface Name—Lets you add a name to a new interface, or displays the name associated with an existing interface.
  - Enable interface—Check this box to activate the interface in privileged mode.
  - Security Level—Displays the security level range for the interface from 0 to 100, with 100 assigned to the inside interface and 0 assigned to the outside interface. Perimeter interfaces can use any number between 1 and 99. Security levels between 0 and 100 for perimeter interfaces are not set by default.
- IP Address area
  - Use PPPoE—Click to obtain an IP address from a PPP over Ethernet (PPPoE) server for the interface.

    The default authentication method for PPPoE is Password Authentication Protocol (PAP). You have the option of configuring Challenge Handshake Authentication Protocol (CHAP) or Microsoft CHAP (MSCHAP) manually.
  - Use DHCP—Click to obtain an IP address from a Dynamic Host Configuration Protocol server so that IP addresses can be reused when hosts no longer need them.

    Note DCHP clients initially have no configured IP address, and must send a broadcast request to obtain an IP address from a DHCP server.

    Obtain default route using DHCP—Check this box to obtain an IP address for the default gateway using DHCP.

    Note DHCP is used by workstations (hosts) to get initial configuration information, such as an IP address, subnet mask, and default gateway upon bootup.
  - Use the following IP address—Click to manually specify an IP address for the interface:

    IP Address—Lets you enter an IP address for an outside interface.

    Subnet Mask—Lets you enter a subnet mask for an outside interface, or alternatively, choose a selection from the drop-down list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > Interfaces

Internet (Outside) VLAN Configuration

Startup Wizard > Internet (Outside) VLAN Configuration

Benefits

The Internet (Outside) VLAN Configuration screen lets you configure your Internet interface by specifying an IP address, or obtaining one from a PPPoE or a DHCP server.

Important Notes

With a full license, the security appliance supports up to five interfaces with a maximum of three interfaces named ‘interface.’ In restricted mode, the security appliance supports up to three interfaces, and in transparent mode, the security appliance supports up to two interfaces.

Once the maximum number of interfaces has been created, or the maximum number of interfaces has already been named, you may not be able to create a new VLAN, and may have to select from an existing VLAN.

Fields

The Internet (Outside) VLAN Configuration screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Select Internet Interface area
  - Choose an interface—Click to choose an interface to configure, then select an interface from the drop-down list.
  - Create new VLAN interface—Click to create a new VLAN interface, then enter the new VLAN number.

    If the maximum number of interfaces has already been created, or the maximum number of interfaces has already been named, you may not be able to create a new VLAN, and may have to select from an existing VLAN. See the Important Notes section for additional information.
- Enable interface—Check this box to activate the interface in privileged mode.
- Interface Name—Lets you specify a name for the interface.
- Security Level—Lets you enter a security level range for the interface from 0 to 100, with 100 assigned to the inside interface and 0 assigned to the outside interface. Perimeter interfaces can use any number between 1 and 99. Security levels between 0 and 100 for perimeter interfaces are not set by default.
- IP Address area

– Use PPPoE—Click to obtain a dynamic IP address from a PPPoE server for an Internet interface.

– Use DHCP—Click to obtain an IP address for the Internet interface from a DHCP server.

Note DCHP clients initially have no configured IP address, and must send a broadcast request to obtain an IP address from a DHCP server.

Obtain default route using DHCP—Check this box to obtain an IP address for the default gateway using DHCP.

Note DHCP is used by workstations (hosts) to get initial configuration information, such as an IP address, subnet mask, and default gateway upon bootup.

– Use the following IP address—Click to specify an IP address for an Internet interface rather than obtaining one from a PPPoE server or DHCP server:

IP Address—Lets you enter an IP address for an Internet interface.

Subnet Mask—Lets you enter a subnet mask for an Internet interface, or alternatively, choose a selection from the drop-down list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > Interfaces

Outside Interface Configuration - PPPoE

Startup Wizard > Outside Interface Configuration - PPPoE

Benefits

The Outside Interface Configuration - PPPoE screen lets you configure your interface by obtaining an IP address from a PPPoE server. The ASA device is the PPPoE on the specified interface.

Before any network layer protocols can be routed, a connection must be opened and negotiated, in this case, using PPPoE authentication.

Fields

The Outside Interface Configuration - PPPoE screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Group Name—Lets you specify the name of the interfaces.

Note You must specify a group name in order to proceed.

- User Authentication area
  - PPPoE Username—Lets you specify the PPPoE username for authentication purposes.
  - PPPoE Password—Lets you specify the PPPoE password for authentication purposes.
  - Confirm PPPoE Password—Lets you confirm the PPPoE password.
- Authentication Method area

  The default authentication method for PPPoE is Password Authentication Protocol (PAP). You have the option of configuring Challenge Handshake Authentication Protocol (CHAP) or Microsoft CHAP (MSCHAP) manually.
  - PAP—Check this to select the Password Authentication Protocol as the authentication method. PAP is the simplest authentication protocol. The username and password are sent unencrypted using this method.
  - CHAP—Check this to select the Challenge Handshake Authentication Protocol method.

    CHAP does not prevent unauthorized access; it merely identifies the remote end. Then, the access server determines whether the user is allowed access.
  - MSCHAP—Check this to select the Microsoft Challenge Handshake Authentication Protocol authentication for PPP connections between a computer using a Microsoft Screens operating system and an access server.
- IP Address area
  - Obtain IP Address using PPPoE—Click to obtain an IP address using a PPPoE server.

    The default authentication method for PPPoE is Password Authentication Protocol (PAP). You have the option of configuring Challenge Handshake Authentication Protocol (CHAP) or Microsoft CHAP (MSCHAP) manually.
  - Specify an IP address—Click to specify an IP address for an interface rather than obtaining one from a PPPoE server:

    IP Address—Lets you enter an IP address for an interface.

    Subnet Mask—Lets you enter a subnet mask for an interface, or alternatively, choose a selection from the drop-down list.
  - Obtain default route using PPPoE—Click to obtain the default route between the PPPoE server and the PPPoE client.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > Interfaces

Internet (Outside) VLAN Configuration - PPPoE

Startup Wizard > Internet (Outside) VAN Configuration - PPPoE

Benefits

The Internet (Outside) VLAN Configuration - PPPoE screen lets you configure your interface by obtaining an IP address from a PPPoE server. The ASA device is the PPPoE on the specified interface.

Before any network layer protocols can be routed, a connection must be opened and negotiated, in this case, using PPPoE authentication.

Fields

The Internet (Outside) VLAN Configuration - PPPoE screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Group Name—Lets you specify the name of the interfaces.

  Note You must specify a group name in order to proceed.

- User Authentication area
  - PPPoE Username—Lets you specify the PPPoE username for authentication purposes.
  - PPPoE Password—Lets you specify the PPPoE password for authentication purposes.
  - Confirm PPPoE Password—Lets you confirm the PPPoE password.
- Authentication Method area

  The default authentication method for PPPoE is Password Authentication Protocol (PAP). You have the option of configuring Challenge Handshake Authentication Protocol (CHAP) or Microsoft CHAP (MSCHAP) manually.
  - PAP—Check this to select the Password Authentication Protocol as the authentication method. PAP is the simplest authentication protocol. The username and password are sent unencrypted using this method.
  - CHAP—Check this to select the Challenge Handshake Authentication Protocol method.

    CHAP does not prevent unauthorized access; it merely identifies the remote end. Then, the access server determines whether the user is allowed access.
  - MSCHAP—Check this to select the Microsoft Challenge Handshake Authentication Protocol authentication for PPP connections between a computer using a Microsoft screens operating system and an access server.
- IP Address area
  - Obtain IP Address using PPPoE—Click to obtain an IP address using a PPPoE server.

    The default authentication method for PPPoE is Password Authentication Protocol (PAP). You have the option of configuring Challenge Handshake Authentication Protocol (CHAP) or Microsoft CHAP (MSCHAP) manually.

– Specify an IP address—Click to specify an IP address for an interface rather than obtaining one from a PPPoE server:

  IP Address—Lets you enter an IP address for an interface.

  Subnet Mask—Lets you enter a subnet mask for an interface, or alternatively, choose a selection from the drop-down list.

– Obtain default route using PPPoE—Click to obtain the default route between the PPPoE server and the PPPoE client.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > Interfaces

Inside Interface Configuration

Startup Wizard > Inside Interface Configuration

Benefits

The Inside Interface Configuration screen lets you configure an inside interface by specifying an IP address, or obtaining one from a PPPoE or a DHCP server.

Note If VLAN is configured, the screen displays a message that in order to make additional changes, you should go to Configuration > Interfaces.

Important Notes

With a full license, the security appliance supports up to five interfaces with a maximum of three interfaces named ‘interface.’ In restricted mode, the security appliance supports up to three interfaces, and in transparent mode, the security appliance supports up to two interfaces.

Once the maximum number of interfaces has been created, or the maximum number of interfaces has already been named, you may not be able to create a new VLAN, and may have to select from an existing VLAN.

Fields

The Inside Interface Configuration screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Select Inside Interface area
  – Choose an interface—Choose an interface to configure from the drop-down list.

  - Create new VLAN interface—Click to create a new inside interface
  - Enable interface—Check this box to activate the interface in privileged mode.
- Interface Name—Lets you specify a name for the interface.
- Security Level—Lets you enter a security level range for the interface from 0 to 100, with 100 assigned to the inside interface and 0 assigned to the outside interface. Perimeter interfaces can use any number between 1 and 99. Security levels between 0 and 100 for perimeter interfaces are not set by default.
- IP Address area
  - Use PPPoE—Click to obtain an IP address from a PPPoE server for an inside interface.
  - Use DHCP—Click to obtain an IP address for the inside interface from a DHCP server.

    Note DCHP clients initially have no configured IP address, and must send a broadcast request to obtain an IP address from a DHCP server.

    Obtain default route using DHCP—Check this box to obtain an IP address for the default gateway using DHCP.

    Note DHCP is used by workstations (hosts) to get initial configuration information, such as an IP address, subnet mask, and default gateway upon bootup.

  - Use the following IP address—Lets you specify an IP address for an inside interface rather than obtaining one from a PPPoE server or DHCP server:

    IP Address—Lets you specify an IP address for an inside interface.

    Subnet Mask—Lets you specify a subnet mask for an inside interface; the list displays a selection of subnet mask IP addresses.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > Interfaces

Business (Inside) VLAN Configuration

Startup Wizard > Business (Inside) VLAN Configuration

Benefits

The Business (Inside) VLAN Configuration screen lets you configure an inside interface by specifying an IP address, or obtaining one from a PPPoE or a DHCP server.

Note If VLAN is configured, the screen displays a message that in order to make additional changes, you should go to Configuration > Interfaces.

Important Notes

With a full license, the security appliance supports up to five interfaces with a maximum of three interfaces named 'interface.' In restricted mode, the security appliance supports up to three interfaces, and in transparent mode, the security appliance supports up to two interfaces.

Once the maximum number of interfaces has been created, or the maximum number of interfaces has already been named, you may not be able to create a new VLAN, and may have to select from an existing VLAN.

Fields

The Business (Inside) VLAN Configuration screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Select Inside Interface area
  - Choose an interface—Choose an interface to configure from the drop-down list.
  - Create new VLAN interface—Click to create a new inside interface
  - Enable interface—Check this box to activate the interface in privileged mode.
- Interface Name—Lets you specify a name for the interface.
- Security Level—Lets you enter a security level range for the interface from 0 to 100, with 100 assigned to the inside interface and 0 assigned to the outside interface. Perimeter interfaces can use any number between 1 and 99. Security levels between 0 and 100 for perimeter interfaces are not set by default.
- IP Address area
  - Use PPPoE—Click to obtain an IP address from a PPPoE server for an inside interface.
  - Use DHCP—Click to obtain an IP address for the inside interface from a DHCP server.

    Note DCHP clients initially have no configured IP address, and must send a broadcast request to obtain an IP address from a DHCP server.

    Obtain default route using DHCP—Check this box to obtain an IP address for the default gateway using DHCP.

    Note DHCP is used by workstations (hosts) to get initial configuration information, such as an IP address, subnet mask, and default gateway upon bootup.

  - Use the following IP address—Lets you specify an IP address for an inside interface rather than obtaining one from a PPPoE server or DHCP server:

    IP Address—Lets you specify an IP address for an inside interface.

Subnet Mask—Lets you specify a subnet mask for an inside interface; the list includes a selection of subnet mask IP addresses.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > Interfaces

DMZ Interface Configuration

Startup Wizard > DMZ Interface Configuration

Benefits

The DMZ Interface Configuration screen lets you configure a work interface.

The security appliance supports up to three fully functional named interfaces; in transparent mode, the security appliance supports up to two interfaces. Typically one interface connects to the outside Internet (known as an Internet zone), another connects to a home network (known as a home zone), and the third interface (known as a work interface), operates similarly to a demilitarized zone (DMZ). A DMZ is a separate network located in the neutral zone between a private (inside) network and a public (outside) network.

Important Notes

With a full license, the security appliance supports up to five interfaces with a maximum of three interfaces named ‘interface.’ In restricted mode, the security appliance supports up to three interfaces, and in transparent mode, the security appliance supports up to two interfaces.

Once the maximum number of interfaces has been created, or the maximum number of interfaces has already been named, you may not be able to create a new VLAN, and may have to select from an existing VLAN.

Fields

The DMZ Interface Configuration screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Select Work Interface area
  - Choose an interface—Choose an interface to configure from the drop-down list.
  - Create new VLAN interface—Check this box to create a new work interface.
  - Enable interface—Check this box to activate the interface in privileged mode.
- Interface Name—Lets you specify a name for the interface.

- Security Level—Lets you enter a security level range for the interface from 0 to 100, with 100 assigned to the inside interface and 0 assigned to the outside interface. Perimeter interfaces can use any number between 1 and 99. Security levels between 0 and 100 for perimeter interfaces are not set by default.
- IP Address area
  - Use PPPoE—Check this box to obtain an IP address from a PPPoE server for a work interface.
  - Use DHCP—Check this box to obtain an IP address for a work interface from a DHCP server.

    Note DCHP clients initially have no configured IP address, and must send a broadcast request to obtain an IP address from a DHCP server.

    Obtain default route using DHCP—Check this box to obtain an IP address for the default gateway using DHCP.

    Note DHCP is used by workstations (hosts) to get initial configuration information, such as an IP address, subnet mask, and default gateway upon bootup.

  - Use the following IP address—Lets you specify an IP address for a work interface rather than obtaining one from a PPPoE server or DHCP server:

    IP Address—Lets you specify an IP address for a work interface.

    Subnet Mask—Lets you specify a subnet mask for a work interface; use the drop-down list to select a subnet mask IP address.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > Interfaces

Home (DMZ) VLAN Configuration

Startup Wizard > Home (DMZ) VLAN Configuration

Benefits

The Home (DMZ) VLAN Configuration screen lets you configure a work interface.

The security appliance supports up to three fully functional named interfaces; in transparent mode, the security appliance supports up to two interfaces. Typically one interface connects to the outside Internet (known as an Internet zone), another connects to a home network (known as a home zone), and the third interface (known as a work interface), operates similarly to a demilitarized zone (DMZ). A DMZ is a separate network located in the neutral zone between a private (inside) network and a public (outside) network.

Important Notes

With a full license, the security appliance supports up to five interfaces with a maximum of three interfaces named 'interface.' In restricted mode, the security appliance supports up to three interfaces, and in transparent mode, the security appliance supports up to two interfaces.

Once the maximum number of interfaces has been created, or the maximum number of interfaces has already been named, you may not be able to create a new VLAN, and may have to select from an existing VLAN.

Fields

The Home (DMZ) VLAN Configuration screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Select Work Interface area
  - Choose an interface—Choose an interface to configure from the drop-down list.
  - Create new VLAN interface—Check this box to create a new work interface.
  - Enable interface—Check this box to activate the interface in privileged mode.
- Interface Name—Lets you specify a name for the interface.
- Security Level—Lets you enter a security level range for the interface from 0 to 100, with 100 assigned to the inside interface and 0 assigned to the outside interface. Perimeter interfaces can use any number between 1 and 99. Security levels between 0 and 100 for perimeter interfaces are not set by default.
- IP Address area
  - Use PPPoE—Check this box to obtain an IP address from a PPPoE server for a work interface.
  - Use DHCP—Check this box to obtain an IP address for a work interface from a DHCP server.

    Note DCHP clients initially have no configured IP address, and must send a broadcast request to obtain an IP address from a DHCP server.

    Obtain default route using DHCP—Check this box to obtain an IP address for the default gateway using DHCP.

    Note DHCP is used by workstations (hosts) to get initial configuration information, such as an IP address, subnet mask, and default gateway upon bootup.

  - Use the following IP address—Lets you specify an IP address for a work interface rather than obtaining one from a PPPoE server or DHCP server:

    IP Address—Lets you specify an IP address for a work interface.

    Subnet Mask—Lets you specify a subnet mask for a work interface; use the drop-down list to display a selection of subnet mask IP addresses.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > Interfaces

Switch Port Allocation

Startup Wizard > Switch Port Allocation

Benefits

The Switch Port Allocation screen lets you allocate switch ports to your outside, inside, and work interface. As VLANs are port-based, you must add the ports to their respective VLANs. By default, all switch ports begin in VLAN1.

Fields

The Switch Port Allocation screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

Allocate Switch Ports to your Outside Interface (vlan*id*) area

- Available Ports—Lets you select a port to add or remove from the available list of ports.
- Allocated Ports—Lets you select a port to add or remove from the allocated list of ports.
- Add—Lets you add a port to the available or allocated list of ports.
- Remove—Lets you remove a port from the available or allocated list of ports.

Allocate Switch Ports to your Inside Interface (vlan*id*) area

- Available Ports—Lets you select a port to add or remove from the available list of ports.
- Allocated Ports—Lets you select a port to add or remove from the allocated list of ports.
- Add—Lets you add a port to the available or allocated list of ports.
- Remove—Lets you remove a port from the available or allocated list of ports.

Allocate Switch Ports to your Work Interface (vlan*id*) area

- Available Ports—Lets you select a port to add or remove from the available list of ports.
- Allocated Ports—Lets you select a port to add or remove from the allocated list of ports.
- Add—Lets you add a port to the available or allocated list of ports.
- Remove—Lets you remove a port from the available or allocated list of ports.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent[1]	Single	Context	System
•	•	•	•	—

1. Work interface is hidden in transparent mode.

For More Information

This feature is available in the main ASDM application screen:

Configuration > Interfaces

General Interface Configuration

Startup Wizard > General Interface Configuration

Benefits

The General Interface Configuration screen lets you enable and restrict traffic between interfaces and between hosts connected to the same interface.

Important Notes

Restricted traffic is not an optional configuration. If you only have a restricted license, you must restrict from one interface to any of the other interfaces. Typically, this is the traffic from the work interface to the inside interface, but any pair can be chosen. The Restrict Traffic area fields are hidden if you have a full license or if the device is in transparent mode.

Fields

The General Interface Configuration screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Enable traffic between two or more interfaces with the same security level—Check this box to enable traffic between two or more interfaces with the same security level.
- Enable traffic between two or more hosts connected to the same interface—Check this box to enable traffic between two or more hosts connected to the same interface.

Restrict traffic area

Note Restricted traffic is not an optional configuration. If you only have a restricted license, you must restrict from one interface to any of the other interfaces. These fields are hidden if you have a full license or if the device is in transparent mode.

- From interface—Lets you restrict traffic from an interface by selecting an interface from the drop-down menu.
- To interface—Lets you restrict traffic to an interface by selecting an interface from the drop-down menu.

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > Interfaces

Static Routes

Startup Wizard >Static Routes

Benefits

The Static Routes screen lets you create static routes that will access networks connected to a router on any interface. To enter a default route, set the IP address and mask to 0.0.0.0, or the shortened form of 0.

If an IP address from one security appliance interface is used as the gateway IP address, the security appliance will ARP the designated IP address in the packet instead of ARPing the gateway IP address.

Leave the Metric at the default of 1 unless you are sure of the number of hops to the gateway router.

Add/Edit Static Routes

Startup Wizard > Static Routes >Add/Edit Static Route

Benefits

The Add/Edit Static Route dialog box lets you add or edit a static route.

Route Monitoring Options

Startup Wizard > Static Routes >Add/Edit Static Route>Route Monitoring Options

Benefits

TheRoute Monitoring Options dialog box lets you set the parameters for monitoring the static route.

Auto Update Server

Startup Wizard > Auto Update Server

Benefits

The Auto Update Server screen allows you to remotely manage the ASA device. This includes automatically updating the ASA configuration, ASA image, and the ASDM image.

Fields

The Auto Update Server screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Enable Auto Update—Check this box to enable communication between the security appliance and an Auto Update Server.
- Server area
  - Server URL—Click the drop-down list to select either the secure http (https) or http to designate the beginning of the URL for the Auto Update server. In the next box, enter the remainder of the IP address for the Auto Update server.
  - Verify server SSL certificate—Check this box to confirm that a SSL certificate is enabled on the Auto Update Server.
- User area
  - Username—Enter the username to log in to the Auto Update server.
  - Password—Enter the password to log in to the Auto Update server.
  - Confirm Password—Enter the password again to confirm it.
- Device Identify area
  - Device ID Type—Click the drop-down list to select the type of ID to uniquely identify the security appliance. Selecting User-defined name enables the Device ID field, where you can specify a unique ID.
  - Device ID—Enter a unique string to use as security appliance ID.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > Interfaces

DHCP Server

Startup Wizard > DHCP Server

Benefits

The DHCP Server screen lets you configure the security appliance as a Dynamic Host Control Protocol (DHCP) server to hosts on the inside interface. You can configure a range of IP addresses in an address pool, then when a host on the inside interface makes a request for an IP address using DHCP, the security appliance assigns it an address from this pool.

Important Notes

- DNS, WINS and other information for interfaces with the lowest security level (inside interfaces) can be set in this screen. To configure the DHCP server for other interfaces, go to the Configuration> Properties > DHCP Services > DHCP Server in the main ASDM screen.
- The number of addresses allowed in the DHCP pool is 256.
- If you configure ASDM to use the DHCP server option, the security appliance uses the inside IP address, adds one address, and configures the address pool based on the number of addresses available according to your feature license and platform. The pool size varies, and might be configured for fewer IP addresses than you are licensed to use to simplify the configuration.

Fields

The DHCP Server screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Enable DHCP server on the inside interface—Check this box to turn on DHCP for the security appliance.
  - DHCP Address Pool area

    Starting IP Address—Enter the starting range of the DHCP server pool in a block of IP addresses from the lowest to highest. The security appliance supports 256 IP addresses.

    Ending IP Address—Enter the ending range of the DHCP server pool in a block of IP addresses from the lowest to highest. The security appliance supports 256 IP addresses.
  - DHCP Parameters area

    Enable auto-configuration—Check this box to allow the wizard to configure the DNS server, WINS server, lease length, and ping timeout.

    DNS Server 1—Enter the IP address of the DNS server to use DNS.

    WINS Server 1—Enter the IP address of the WINS (screens Internet Naming Service) server to use DNS.

    DNS Server 2—Enter the IP address of the alternate DNS server to use DNS.

    WINS Server 2—Enter the IP address of the alternate WINS server to use DNS.

    Lease Length (secs)—Enter the amount of time (in seconds) the client can use its allocated IP address before the lease expires. The default value is 3600 seconds (1 hour).

    Ping Timeout—Enter the parameters for the ping timeout value in milliseconds.

    Domain Name—Enter the domain name of the DNS server to use DNS.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > Properties > DHCP Services > DHCP Server

Address Translation (NAT/PAT)

Startup Wizard > Address Translation (NAT/PAT)

Note This feature is not available in transparent mode.

Benefits

The Address Translation (NAT/PAT) screen lets you configure NAT and PAT on your security appliance.

PAT lets you set up a single IP address for use as the global address. With PAT, you can set multiple outbound sessions to appear as if they originate from a single IP address. When enabled, the security appliance chooses a unique port number from the PAT IP address for each outbound translation slot. This feature is useful in smaller installations where there are not enough unique IP addresses for all outbound connections. An IP address that you specify for a port address cannot be used in another global address pool. PAT lets up to 65,535 hosts start connections through a single outside IP address.

If you decide to use NAT, enter an address range to use for translating addresses on the inside interface to addresses on the outside interface. The global addresses in the pool provide an IP address for each outbound connection, and for those inbound connections resulting from outbound connections.

Important Notes

If you use NAT, the range of IP addresses required on this screen creates a pool of addresses that is used outbound on the security appliance. If you have been assigned a range of Internet-registered, global IP addresses by your ISP, enter them here.

The following are limitations when using the PAT address configuration:

- Does not work with caching name servers.
- You may need to enable the corresponding inspection engine to pass multimedia application protocols through the security appliance.
- Does not work with the established command.
- When in use with a passive FTP, use the **Inspect protocol ftp strict** command statement with an **access-list** command statement to permit outbound FTP traffic.
- A DNS server on a higher level security interface, needing to get updates from a root name server on the outside interface, cannot use PAT.

Fields

The Address Translation (NAT/PAT) screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Enable traffic through the firewall without translation—Click to allow traffic through the firewall without translation.

  NAT is a one-to-one address translation; PAT is a many (inside the firewall)-to-one translation.

- Use Network Address Translation (NAT)—Select to enable NAT and a range of IP addresses to be used for translation.
  - Starting Global IP Address—Enter the first IP address in a range of IP addresses to be used for translation.

  - Ending Global IP Address—Enter the last IP address in a range of IP addresses to be used for translation.
  - Subnet Mask (optional)—Specify the subnet mask for the range of IP addresses to be used for translation.

- Use Port Address Translation (PAT)—Select to enable PAT. You must choose one of the following if you select this option.

Note IPSec with PAT may not work properly because the outside tunnel endpoint device cannot handle multiple tunnels from one IP address.

  - Use the IP address on the outside interface—The security appliance uses the IP address of the outside interface for PAT.
  - Specify an IP address—Specify an IP address to be used for PAT.

    IP Address—Lets you enter an IP address for the outside interface for PAT.

    Subnet Mask (optional)—Lets you enter a subnet mask for the outside interface for PAT, or click the down arrow to select a subnet mask.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > NAT

Administrative Access

Startup Wizard > Administrative Access

Benefits

The Administrative Access screen lets you configure management access on the security appliance. ASDM automatically lists the interfaces available for configuration, and in this screen you can set the IP address, interface name, and security level to make each interface unique.

Note This screen allows configuration of management access to interfaces already configured in other places. User cannot change such things as the IP address and the name of the interface in this screen.

Fields

The Administrative Access screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Type—Specifies whether the host or network is accessing security appliance via HTTPS/ASDM, SSH, or Telnet.
- Interface—Displays the host or network name.
- IP Address—Displays the IP address of the host or network.
- Mask—Displays the subnet mask of the host or network.
- Add—Lets you choose access type, an interface, then specify the IP address and netmask of the host/network that will be allowed to connect to that interface for management purposes only.
- Edit—Lets you edit an interface.
- Delete—Lets you delete an interface.

Add/Edit Administrative Access Entry

Startup Wizard > Add/Edit Administrative Access Entry

Benefits

The Add/Edit Administrative Access Entry dialog box let you configure the hosts.

You must use one the following types of preconfigured connections for the Command Line Interface console sessions:

- Telnet protocol—A network connection using the Telnet protocol.
- ASDM/HTTPS protocol—A network connection using the HTTPS (HTTP over SSL) protocol for **Tools > Command Line Interface**.

Note ASDM uses HTTPS for all communication with the security appliance.

- Secure Shell (SSH) protocol—A network connection using the Secure Shell (SSH) protocol.

Before configuring your security appliance from the ASDM Command Line Interface tool, we recommend that you review the security appliance Technical Documentation. See also Password, Authentication.

For more information about the Command Line Interface commands used by each ASDM screen, see **Command Line Interface Commands Used by ASDM screens Help > About the security appliance** that will display, among other useful things, which user last changed the configuration.

Fields

The Add/Edit Administrative Access Entry screen displays the OK, Cancel, and Help buttons, in addition to the following:

- Access Type—Select one of the following types of preconfigured connections for the Command Line Interface console sessions from the drop-down menu: ASDM/HTTP, SSH, or Telnet.
- Interface Name—Lets you select from a list of predetermined interfaces.
- IP Address—Lets you specify an IP address for the interface.

- Subnet Mask—Lets you specify a subnet mask for the interface from a selection of subnet mask IP addresses.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > Properties > Device Access > HTTPS/ASDM

Configuration > Properties > Device Access > Telnet

Configuration > Properties > Device Access > SSH

Configuration > Properties > History Metrics

Easy VPN Remote Configuration

Startup Wizard >Easy VPN Remote Configuration

Benefits

Companies with multiple sites can establish secure communications and resource sharing among them by deploying a Cisco Easy VPN solution that consists of an Easy VPN Server at its main site and Easy VPN remote devices at remote offices. Using an Easy VPN solution simplifies the deployment and management of a Virtual Private Network in the following ways:

- Hosts at remote sites no longer have to run VPN client software.
- Security policies reside on a central server and are pushed to the remote devices when a VPN connection is established.
- Few configuration parameters need to be set locally.

When used as an Easy VPN remote device, the security appliance can also be configured to perform basic firewall services, such as protecting a web server on a DMZ from unauthorized access. However, if the security appliance is configured to function as an Easy VPN remote device, it cannot establish other types of tunnels. For example, the security appliance cannot function simultaneously as both an Easy VPN remote device and as one end of a standard peer-to-peer VPN deployment.

The Easy VPN Remote Configuration screen lets you form a secure VPN tunnel between the security appliance and a remotely located Cisco VPN 3000 Concentrator, Cisco IOS-based router, or security appliance that is acting as an Easy VPN server. The security appliance itself acts as an Easy VPN remote device to enable deployment of VPNs to remote locations via the devices listed above.

There are two modes of operation when using the security appliance as an Easy VPN remote device:

- Client Mode
- Network Extension Mode

When configured in Easy VPN Client Mode, the security appliance does not expose the IP addresses of clients on its inside network. Instead, it uses NAT (Network Address Translation) to translate the IP addresses on the private network to a single, assigned IP address. When thesecurity appliance is configured in Client Mode, you cannot ping or access any device from outside the private network.

When configured in Easy VPN Network Extension Mode, the security appliance does not protect the IP addresses of local hosts by substituting a assigned IP address. Therefore, hosts on the other side of the VPN connection can communicate directly with hosts on the local network.

Use the following guidelines when deciding whether to configure the security appliance in Easy VPN Client or Network Extension Mode:

Use Client Mode if:

- You want VPN connections to be initiated by client traffic
- You want the IP addresses of local hosts to be hidden from remote networks
- You are using DHCP on the ASA 5505 to provide IP addresses to local hosts.

Use Network Extension Mode if:

- You want VPN connections to remain open even when not required for transmitting traffic.
- You want remote hosts to be able to communicate directly with hosts on the local network.
- Hosts on the local network have static IP addresses.

Important Notes

- ASA supports a maximum of 11 Easy VPN Servers: one primary and up to 10 secondary.
- In Easy VPN Client Mode, you use a DHCP server to generate dynamic IP addresses for hosts on the inside network.

  To use Easy VPN Client Mode, you must enable the DHCP server on the inside interface.
- Before you can connect the ASA Easy VPN remote device to the Easy VPN Server, you must establish network connectivity between both devices through your Internet service provider (ISP).

  After connecting your ASA to the DSL or cable modem, you should follow the instructions provided by your ISP to complete the network connection. Basically, there are three methods of obtaining an IP address when establishing connectivity to your ISP:

  – PPPoE client
  – DHCP client
  – Static IP address configuration

  The Easy VPN Server controls the policy enforced on the ASA Easy VPN remote device. However, to establish the initial connection to the Easy VPN Server, you must complete some configuration locally.

  You can perform this configuration by following the steps in this Wizard or by using the command-line interface.

Fields

The Easy VPN Remote Configuration screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Enable Easy VPN remote—Check this box to enable the ASA to act as an Easy VPN remote device. Enabling the ASA to act as an Easy VPN Remote allows you to choose the networks from which your ASA can be remotely managed. If you do not enable this feature, any host that has access to the ASA outside interface through a VPN tunnel can manage it remotely.

- Mode area
  - Client Mode—Click if you are using a DHCP server to generate dynamic IP addresses for hosts on your inside network.

    Client Mode enables VPN connections by traffic, allowing resources to be only used on demand. The ASA applies Network Address Translation (NAT) to all IP addresses of clients connected to the inside (higher security) interface of the ASA.

    Note To use Client Mode, you must enable the DHCP server on the inside interface.

  - Network extension—Click if hosts on your inside network have static IP addresses.

    In Network Extension Mode, IP addresses of clients on the inside interface are received without change at the Easy VPN Server, and VPN connections are kept open even when not required for transmitting traffic. This option does not apply NAT to any IP addresses of clients on the inside (higher security) interface of the ASA.
- Group Settings area
  - Use X.509 Certificate—Click to use X.509 certificates to allow IPSec Main Mode. Use the drop-down list to select a trustpoint or to enter a trustpoint.
  - Use group password—Lets you enter a password for a group of users.

    Group Name—Lets you enter a name for the user group.

    Password—Lets you enter a password for the user group.

    Confirm password—Requires that you confirm the password.
- User Settings area
  - Username—Lets you enter a username for your settings.
  - Password—Lets you enter a password for your settings.
  - Confirm Password—Requires that you confirm the password for your settings.
- Easy VPN Server area—Using the ASA as an Easy VPN Server lets you configure your VPN policy in a single location on the ASA, and then push this configuration to multiple Easy VPN remote devices.
  - Primary server—Lets you enter the IP address of the primary Easy VPN Server.
  - Secondary server—Lets you enter the IP address of a secondary Easy VPN Server.

    Note ASA supports a maximum of 11 Easy VPN Servers (one primary and up to 10 secondary).

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > Interfaces

Management IP Address Configuration

Startup Wizard > Management IP Address Configuration

Note This feature is available only in transparent mode.

Benefits

The Management IP Address Configuration screen lets you configure the management IP address of the host for this context.

Fields

The Management IP Address Configuration screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Management IP Address—The IP address of the host that can access this context for management purposes using ASDM or a session protocol.
- Subnet Mask—Subnet mask for the Management IP address.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	—	—	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > Properties > Management IP

Other Interfaces Configuration

Startup Wizard > Other Interfaces Configuration

Benefits

The Other Interfaces Configuration screen lets you configure the remaining interfaces. You highlight a listed interface, select the Edit button, and configure it from the Edit screen.

Fields

The Other Interfaces Configuration screen displays the Back, Next, Finish, Cancel, and Help buttons, in addition to the following:

- Interface—Displays the network interface on which the original host or network resides.
- Name—Displays the name of the interface being edited.
- Security Level—Displays the security level range for the interface from 0 to 100, with 100 assigned to the inside interface and 0 assigned to the outside interface. Perimeter interfaces can use any number between 1 and 99. Security levels between 0 and 100 for perimeter interfaces are not set by default.
- Enable traffic between two or more interfaces with same security levels—Check this box if you assign the same security level to two or more interfaces, and want to enable traffic between the interfaces.
- Enable traffic between two or more hosts connected to the same interface—Check this box if you have an interface between two or more hosts and want to enable traffic between them.
- Edit—Click **Edit** to configure the interface in the Edit Interface dialog box.

Edit Interface

Startup Wizard > Other Interfaces Configuration >Edit Interface

Benefits

Use the Edit Interfaces to edit exisiting interfaces.

Fields

The Edit Interface dialog box displays the OK, Cancel, and Help buttons, in addition to the following:

- Interface—Displays the name of the selected interface to edit.
- Interface Name—Displays the name of the selected interface, and lets you change the name of the interface.
- Security Level—Displays the security level of the selected interface, or lets you select a security level for the interface. Either 0 for the outside network or 100 for the inside network. Perimeter interfaces can use any number between 1 and 99. Security levels between 0 and 100 for perimeter interfaces are not set by default. If you change the security level of the interface to a lower level, a caution warning appears.
- IP Address area
  - Use PPPoE—Check this box to use PPPoE to provide an authenticated method of assigning an IP address to an outside interface. PPPoE provides a standard method of employing the authentication methods of the Point-to-Point Protocol (PPP) over an Ethernet network.

    Note Because PPPoE is permitted on multiple interfaces, each instance of the PPPoE client may require different authentication levels with different usernames and passwords.

  - Use DHCP—Check this box to use ASA as a DHCP server to provide network configuration parameters, including dynamically assigned IP addresses, to DHCP clients.
  - Uses the following IP address—Check this box to input a specific IP address for an interface.

IP Address—Lets you edit the IP address of the interface.

Subnet Mask—Lets you edit the subnet mask by entering a new address or selecting an existing IP address from the drop-down list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

This feature is available in the main ASDM application screen:

Configuration > Interfaces

Startup Wizard Summary

Startup Wizard > Startup Wizard Summary

Benefits

The Startup Wizard Summary screen lets you submit all of the settings you made to the security appliance.

- If you would like to change any of the settings you made, click **Back**.
- If you started the Startup Wizard directly from a browser, when you click **Finish**, the configuration created by the wizard is sent to the security appliance and saved to Flash memory.
- If you ran the Startup Wizard from within ASDM, you must explicitly save the configuration to Flash memory just like any other configuration changes.

Fields

The Startup Wizard Summary screen displays the Back, Finish, Cancel and Help buttons.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Configuring Interfaces

This chapter describes how to configure each interface and subinterface for a name, security level, and IP address. In multiple context mode, you can configure hardware properties and create subinterfaces in the system execution space, while you configure the IP address, name, and security level in each context.

Note To configure interfaces for the ASA 5505 adaptive security appliance, see Chapter 5, "Configuring Switch Ports and VLAN Interfaces for the Cisco ASA 5505 Adaptive Security Appliance."

This chapter includes the following sections:

Security Level Overview

Each interface must have a security level from 0 (lowest) to 100 (highest). For example, you should assign your most secure network, such as the inside host network, to level 100. While the outside network connected to the Internet can be level 0. Other networks, such as DMZs can be in between. You can assign interfaces to the same security level.

The level controls the following behavior:

- Network access—By default, there is an implicit permit from a higher security interface to a lower security interface (outbound). Hosts on the higher security interface can access any host on a lower security interface. You can limit access by applying an access list to the interface.

  For same security interfaces, there is an implicit permit for interfaces to access other interfaces on the same security level or lower.

- Inspection engines—Some application inspection engines are dependent on the security level. For same security interfaces, inspection engines apply to traffic in either direction.
  - NetBIOS inspection engine—Applied only for outbound connections.
  - SQL*Net inspection engine—If a control connection for the SQL*Net (formerly OraServ) port exists between a pair of hosts, then only an inbound data connection is permitted through the security appliance.

- Filtering—HTTP(S) and FTP filtering applies only for outbound connections (from a higher level to a lower level).

  For same security interfaces, you can filter traffic in either direction.

- NAT control—When you enable NAT control, you must configure NAT for hosts on a higher security interface (inside) when they access hosts on a lower security interface (outside).

  Without NAT control, or for same security interfaces, you can choose to use NAT between any interface, or you can choose not to use NAT. Keep in mind that configuring NAT for an outside interface might require a special keyword.

- **established** command—This command allows return connections from a lower security host to a higher security host if there is already an established connection from the higher level host to the lower level host.

  For same security interfaces, you can configure **established** commands for both directions.

Configuring the Interfaces

By default, all physical interfaces are shut down. You must enable the physical interface before any traffic can pass through an enabled subinterface. For multiple context mode, if you allocate a physical interface or subinterface to a context, the interfaces are enabled by default in the context. However, before traffic can pass through the context interface, you must also enable the interface in the system configuration. If you shut down an interface in the system execution space, then that interface is down in all contexts that share it.

Note If you are using failover, do not use this procedure to name interfaces that you are reserving for failover and Stateful Failover communications. See Chapter 12, "Failover." to configure the failover and state links.

For multiple context mode, follow these guidelines:

- Configure the context interfaces from within each context.
- You can only configure context interfaces that you already assigned to the context in the system configuration.
- The system configuration only lets you configure Ethernet settings and VLANs. The exception is for failover interfaces; do not configure failover interfaces with this procedure. See the Failover chapter for more information.

This section includes the following topics:

Interfaces (System)

System > Configuration > Interfaces

The Interfaces pane displays configured interfaces and subinterfaces. Before you can assign an interface to a security context (see the "Configuring Security Contexts" section on page 7-16), define the interface in this pane. Although the system configuration does not include any networking parameters for these interfaces, the system controls the allocation of interfaces to security contexts.

Fields

- Interface—Displays the interface ID. All physical interfaces are listed automatically. Subinterfaces are indicated by the interface ID followed by **.***n*, where *n* is the subinterface number.

  If you use failover, you need to assign a dedicated physical interface as the failover link and an optional interface for Stateful Failover on the Failover: Setup tab. (You can use the same interface for failover and state traffic, but we recommend separate interfaces). To ensure that you can use an interface for failover, do not configure an interface name in the Interfaces pane. Other settings, including the IP address, are ignored; you set all relevant parameters in the Failover: Setup tab. You can use a subinterface for failover as long as you do not set a name for the physical interface or the subinterface. After you assign an interface as the failover link or state link, you cannot edit or delete the interface in this pane. The only exception is if you set a physical interface to be the state link, then you can configure the speed and duplex.

- Enabled—Indicates if the interface is enabled, Yes or No.

  By default, all physical interfaces are shut down. You must enable the physical interface before any traffic can pass through an enabled subinterface. For multiple context mode, if you allocate a physical interface or subinterface to a context, the interfaces are enabled by default in the context. However, before traffic can pass through the context interface, you must also enable the interface in the system configuration. If you shut down an interface in the system execution space, then that interface is down in all contexts that share it.

- VLAN—Shows the VLAN assigned to a subinterface. Physical interfaces show “native,” meaning that the physical interface is untagged.

- Description—Displays a description. In the case of a failover or state link, the description is fixed as “LAN Failover Interface,” “STATE Failover Interface,” or “LAN/STATE Failover Interface,” for example. You cannot edit this description.

- Add—Adds a subinterface.

- Edit—Edits the selected interface. If you assign an interface as the failover link or state link (see the Failover: Setup tab), you cannot edit the interface in this pane. The only exception is if you set a physical interface to be the state link, then you can configure the speed and duplex.

- Delete—Deletes the selected subinterface. You cannot delete physical interfaces or allocated interfaces in a context. If you assign an interface as the failover link or state link (see the Failover: Setup tab), you cannot delete the interface in this pane.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

Add/Edit Interface

System > Configuration > Interfaces > Add/Edit Interface

The Add Interface dialog box lets you add a subinterface. The Edit Interface dialog box lets you edit an interface or subinterface.

If you intend to use a physical interface for failover, do not configure the interface in this dialog box; instead, use the Failover: Setup tab. In particular, do not set the interface name, as this parameter disqualifies the interface from being used as the failover link; other parameters are ignored.

After you assign the interface as the failover link or state link, you cannot edit or delete the interface from the Interfaces pane. The only exception is if you set a physical interface to be the state link, then you can configure the speed and duplex.

Fields

- Hardware Port—When you add a subinterface, you can choose any enabled physical interface to which you want to add a subinterface. If you do not see an interface ID, be sure that the interface is enabled.
- Configure Hardware Properties—For a physical interface, opens the Hardware Properties dialog box so you can set the speed and duplex.
- Enable Interface—Enables this interface to pass traffic.

  By default, all physical interfaces are shut down. You must enable the physical interface before any traffic can pass through an enabled subinterface. For multiple context mode, if you allocate a physical interface or subinterface to a context, the interfaces are enabled by default in the context. However, before traffic can pass through the context interface, you must also enable the interface in the system configuration. If you shut down an interface in the system execution space, then that interface is down in all contexts that share it.
- VLAN ID—For a subinterface, sets the VLAN ID, between 1 and 4095. Some VLAN IDs might be reserved on connected switches, so check the switch documentation for more information. For multiple context mode, you can only set the VLAN in the system configuration.
- Sub-interface ID—Sets the subinterface ID as an integer between 1 and 4294967293. The number of subinterfaces allowed depends on your platform. You cannot change the ID after you set it.
- Description—Sets an optional description up to 240 characters on a single line, without carriage returns. The system description is independent of the context description. In the case of a failover or state link, the description is fixed as “LAN Failover Interface,” “STATE Failover Interface,” or “LAN/STATE Failover Interface,” for example. You cannot edit this description. The fixed description overwrites any description you enter here if you make this interface a failover or state link.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

Hardware Properties

System > Configuration > Interfaces > Edit Interface > Hardware Properties

The Hardware Properties dialog box lets you set the speed and duplex of physical interfaces.

Fields

- Hardware Port—*Display only*. Displays the interface ID.
- Media Type—Sets the media type to RJ45 or SFP. The default is RJ45.
- Duplex—Lists the duplex options for the interface, including Full, Half, or Auto, depending on the interface type.
- Speed—Lists the speed options for the interface. The speeds available depend on the interface type. For SFP interfaces, which are always 1000 Mbps, and you can set the speed to Negotiate or Nonegotiate. Negotiate (the default) enables link negotiation, which exchanges flow-control parameters and remote fault information. Nonegotiate does not negotiate link parameters. For RJ-45 interfaces on the ASA 5500 series adaptive security appliance, the default auto-negotiation setting also includes the Auto-MDI/MDIX feature. Auto-MDI/MDIX eliminates the need for crossover cabling by performing an internal crossover when a straight cable is detected during the auto-negotiation phase. Either the speed or duplex must be set to auto-negotiate to enable Auto-MDI/MDIX for the interface. If you explicitly set both the speed and duplex to a fixed value, thus disabling auto-negotiation for both settings, then Auto-MDI/MDIX is also disabled.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

Interfaces (Single Mode and Context)

Configuration > Interfaces

The Interfaces pane displays configured interfaces and subinterfaces. You can add or delete subinterfaces (single mode only), and also enable communication between interfaces on the same security level or enable traffic to enter and exit the same interface.

Transparent firewall mode allows only two interfaces to pass through traffic; however, if your platform includes a dedicated management interface, Management 0/0, you can use it (either the physical interface or a subinterface) as a third interface for management traffic.

Benefits

This pane lets you enable communication between interfaces on the same security level.

By default, interfaces on the same security level cannot communicate with each other. Allowing communication between same security interfaces provides the following benefits:

- You can configure more than 101 communicating interfaces.

  If you use different levels for each interface and do not assign any interfaces to the same security level, you can configure only one interface per level (0 to 100).
- You want traffic to flow freely between all same security interfaces without access lists.

Fields

- Interface—Displays the interface ID. All physical interfaces are listed automatically. Subinterfaces are indicated by the interface ID followed by **.***n*, where *n* is the subinterface number.

  If you use failover, you need to assign a dedicated physical interface as the failover link and an optional interface for Stateful Failover on the Failover: Setup tab. (You can use the same interface for failover and state traffic, but we recommend separate interfaces). To ensure that you can use an interface for failover, do not configure an interface name in the Interfaces pane. Other settings, including the IP address, are ignored; you set all relevant parameters in the Failover: Setup tab. You can use a subinterface for failover as long as you do not set a name for the physical interface or the subinterface. After you assign an interface as the failover link or state link, you cannot edit or delete the interface in this pane. The only exception is if you set a physical interface to be the state link, then you can configure the speed and duplex.

  For multiple context mode, the physical interfaces are listed only in the system configuration. When you allocate interfaces to a context, each allocated interface is listed automatically in the context.

- Name—Displays the interface name.
- Enabled—Indicates if the interface is enabled, Yes or No. By default, all physical interfaces are shut down. You must enable the physical interface before any traffic can pass through an enabled subinterface. For multiple context mode, if you allocate a physical interface or subinterface to a context, the interfaces are enabled by default in the context. However, before traffic can pass through the context interface, you must also enable the interface in the system configuration. If you shut down an interface in the system execution space, then that interface is down in all contexts that share it.
- Security Level—Displays the interface security level between 0 and 100. By default, the security level is 0.
- IP Address—Displays the IP address, or in transparent mode, the word “native.” Transparent mode interfaces do not use IP addresses. To set the IP address for the context or the security appliance, see the Management IP pane.
- Subnet Mask—For routed mode only. Displays the subnet mask.
- Management Only—Indicates if the interface allows traffic to the security appliance or for management purposes only.
- MTU—Displays the MTU. By default, the MTU is 1500.
- Active MAC Address—Shows the active MAC address, if you assigned one manually on the Add/Edit Interface > Advanced tab.
- Standby MAC Address—Shows the standby MAC address (for failover), if you assigned one manually.
- Description—Displays a description. In the case of a failover or state link, the description is fixed as “LAN Failover Interface,” “STATE Failover Interface,” or “LAN/STATE Failover Interface,” for example. You cannot edit this description.
- Add—Adds a subinterface.
- Edit—Edits the selected interface. If you assign an interface as the failover link or state link (see the Failover: Setup tab), you cannot edit the interface in this pane. The only exception is if you set a physical interface to be the state link, then you can configure the speed and duplex.
- Delete—Deletes the selected subinterface. You cannot delete physical interfaces or allocated interfaces in a context. If you assign an interface as the failover link or state link (see the Failover: Setup tab), you cannot delete the interface in this pane.

- Enable traffic between two or more interfaces which are configured with same security levels—Enables communication between interfaces on the same security level. If you enable same security interface communication, you can still configure interfaces at different security levels as usual.
- Enable traffic between two or more hosts connected to the same interface—Enables traffic to enter and exit the same interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System[1]
•	•	•	•	—

1. For the system Interfaces pane, see the system Interfaces (System) pane.

Add/Edit Interface > General

Configuration > Interfaces > Add/Edit Interface > General

The Add Interface > General tab lets you add a subinterface. The Edit Interface > General tab lets you edit an interface or subinterface. In multiple context mode, you can only add interfaces in the system configuration. See the "Configuring Security Contexts" section on page 7-16 to assign interfaces to contexts.

If you intend to use a physical interface for failover, do not configure the interface in this dialog box; instead, use the Failover: Setup tab. In particular, do not set the interface name, as this parameter disqualifies the interface from being used as the failover link; other parameters are ignored.

After you assign the interface as the failover link or state link, you cannot edit or delete the interface from the Interfaces pane. The only exception is if you set a physical interface to be the state link, then you can configure the speed and duplex.

Fields

- Hardware Port—When you add a subinterface, you can choose any enabled physical interface to which you want to add a subinterface. If you do not see an interface ID, be sure that the interface is enabled.
- Configure Hardware Properties—For a physical interface, opens the Hardware Properties dialog box so that you can set the speed and duplex, and for some interfaces, the media type. For multiple context mode, you can only set physical properties in the system configuration.
- Enable Interface—Enables this interface to pass traffic. In addition to this setting, you need to set an IP address (for routed mode) and a name before traffic can pass according to your security policy. By default, all physical interfaces are shut down. You must enable the physical interface before any traffic can pass through an enabled subinterface. For multiple context mode, if you allocate a physical interface or subinterface to a context, the interfaces are enabled by default in the context. However, before traffic can pass through the context interface, you must also enable the interface in the system configuration. If you shut down an interface in the system execution space, then that interface is down in all contexts that share it.
- Dedicate this interface to management only—Sets the interface to accept traffic to the security appliance only, and not through traffic.

- VLAN ID—For a subinterface, sets the VLAN ID, between 1 and 4095. Some VLAN IDs might be reserved on connected switches, so check the switch documentation for more information. For multiple context mode, you can only set the VLAN in the system configuration.
- Sub-interface ID—Sets the subinterface ID as an integer between 1 and 4294967293. The number of subinterfaces allowed depends on your platform. You cannot change the ID after you set it.
- Interface Name—Sets an interface name up to 48 characters in length.
- Security Level—Sets the security level between 0 (lowest) and 100 (highest).The security appliance lets traffic flow freely from an inside network to an outside network (lower security level). Many other security features are affected by the relative security level of two interfaces.
- IP Address—For routed mode only. For multiple context mode, set the IP address in the context configuration.
  - Use Static IP—Manually sets the IP address.

    IP address—Sets the IP address.

    Subnet Mask—Sets the subnet mask.
  - Obtain Address via DHCP—Dynamically sets the IP address using DHCP.

    Obtain Default Route Using DHCP—Obtains a default route from the DHCP server so that you do not need to configure a default static route.

    Renew DHCP Lease—Renews the DHCP lease.

    Retry Count—Sets the number of times between 4 and 16 that the security appliance resends a DHCP request if it does not receive a reply after the first attempt. The total number of attempts is the retry count plus the first attempt. For example, if you set the retry count to 4, the security appliance sends up to 5 DHCP requests.

    DHCP Learned Route Metric—Assigns an administrative distance to the learned route. Valid values are from 1 to 255. If this field is left blank, the administrtive distance for the learned routes is 1.

    Enable tracking—Check this checkbox to enable route tracking for DHCP-learned routes.

    Note Route tracking is only available in single, routed mode.

    Track ID—A unique identifier for the route tracking process. Valid values are from 1 to 500.

    Track IP Address—Enter the IP address of the target being tracked. Typically, this would be the IP address of the next hop gateway for the route, but it could be any network object available off of that interface.

    SLA ID—A unique identifier for the SLA monitoring process. Valid values are from 1 to 2147483647.

    Monitor Options—Click this button to open the Route Monitoring Options dialog box. In the Route Monitoring Options dialog box you can configure the parameters of the tracked object monitoring process.
  - Use PPPoE—Dynamically sets the IP address using PPPoE.

    Note PPPoE is not supported with failover, or in Multiple context mode and Transparent mode. PPPoE is only supported in Single-Routed mode without failover.

    Group Name—Specify a group name.

PPPoE Username—Specify the username provided by your ISP.

PPPoE Password—Specify the password provided by your ISP.

Confirm Password—Specify the password provided by your ISP.

PPP Authentication—Select either PAP, CHAP, or MSCHAP. PAP passes cleartext username and password during authentication and is not secure. With CHAP, the client returns the encrypted [challenge plus password], with a cleartext username in response to the server challenge. CHAP is more secure than PAP, but it does not encrypt data. MSCHAP is similar to CHAP but is more secure because the server stores and compares only encrypted passwords rather than cleartext passwords as in CHAP. MSCHAP also generates a key for data encryption by MPPE.

Store Username and Password in Local Flash—Stores the username and password in a special location of NVRAM on the security appliance. If an Auto Update Server sends a clear config command to the security appliance and the connection is then interrupted, the security appliance can read the username and password from NVRAM and re-authenticate to the Access Concentrator.

IP Address and Route Settings—displays the PPPoE IP Address and Route Settings dialog where you can choose addressing and tracking options.

- Description—Sets an optional description up to 240 characters on a single line, without carriage returns. For multiple context mode, the system description is independent of the context description. In the case of a failover or state link, the description is fixed as "LAN Failover Interface," "STATE Failover Interface," or "LAN/STATE Failover Interface," for example. You cannot edit this description. The fixed description overwrites any description you enter here if you make this interface a failover or state link.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System[1]
•	•	•	•	—

1. For the system Add/Edit Interfaces dialog box, see the system Add/Edit Interface dialog box.

Add/Edit Interface > Advanced

Configuration > Interfaces > Add/Edit Interface > Advanced

The Add/Edit Interface > Advanced tab lets you set the MTU and MAC address of the interface.

Fields

- MTU—Sets the MTU from 300 to 65,535 bytes. The default is 1500 bytes. For multiple context mode, set the MTU in the context configuration.
- Mac Address Cloning—Manually assigns MAC addresses.

  By default, the physical interface uses the burned-in MAC address, and all subinterfaces of a physical interface use the same burned-in MAC address.

In multiple context mode, if you share an interface between contexts, you can assign a unique MAC address to the interface in each context. This feature lets the security appliance easily classify packets into the appropriate context. Using a shared interface without unique MAC addresses is possible, but has some limitations. See the "How the Security Appliance Classifies Packets" section on page 7-2 for more information. You can assign each MAC address manually, or you can automatically generate MAC addresses for shared interfaces in contexts. See the "Security Contexts" section on page 7-16 to automatically generate MAC addresses. If you automatically generate MAC addresses, you can use this option to override the generated address.

For single context mode, or for interfaces that are not shared in multiple context mode, you might want to assign unique MAC addresses to subinterfaces. For example, your service provider might perform access control based on the MAC address.

- Active Mac Address—Assigns a MAC address to the interface in H.H.H format, where H is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE.
- Standby Mac Address—For use with failover, set the Standby Mac Address. If the active unit fails over and the standby unit becomes active, the new active unit starts using the active MAC addresses to minimize network disruption, while the old active unit uses the standby address.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System[1]
•	•	•	•	—

1. For the system Add/Edit Interfaces dialog box, see the system Add/Edit Interface dialog box.

PPPoE IP Address and Route Settings

Configuration > Interfaces > Add/Edit Interface > General > PPPoE IP Address and Route Settings

The PPPoE IP Address and Route Settings dialog lets you choose addressing and tracking options for PPPoE connections.

Fields

- IP Address area—Lets you choose between Obtaining an IP address using PPP or specifying an IP address, and contains the following fields:
  - Obtain IP Address using PPP—Select to enable the security appliance to use PPP to get an IP address.
  - Specify an IP Address—Specify an IP address and mask for the security appliance to use instead of negotiating with the PPPoE server to assign an address dynamically.
- Route Settings Area—Lets you configure route and tracking settings and contains the following fields:

- Obtain default route using PPPoE—Sets the default routes when the PPPoE client has not yet established a connection. When using this option, you cannot have a statically defined route in the configuration.

  PPPoE learned route metric—Assigns an administrative distance to the learned route. Valid values are from 1 to 255. If this field is left blank, the administrtive distance for the learned routes is 1.

- Enable tracking—Check this checkbox to enable route tracking for PPPoE-learned routes.

  Note Route tracking is only available in single, routed mode.

- Primary Track—Select this option to configure the primary PPPoE route tracking.
- Track ID—A unique identifier for the route tracking process. Valid values are from 1 to 500.
- Track IP Address—Enter the IP address of the target being tracked. Typically, this would be the IP address of the next hop gateway for the route, but it could be any network object available off of that interface.
- SLA ID—A unique identifier for the SLA monitoring process. Valid values are from 1 to 2147483647.
- Monitor Options—Click this button to open the Route Monitoring Options dialog box. In the Route Monitoring Options dialog box you can configure the parameters of the tracked object monitoring process.
- Secondary Track—Select this option to configure the secondary PPPoE route tracking.
- Secondary Track ID—A unique identifier for the route tracking process. Valid values are from 1 to 500.

Hardware Properties

Configuration > Interfaces > Edit Interface > Hardware Properties

The Hardware Properties dialog box lets you set the speed and duplex of physical interfaces, and for an interface SSM, the media type. In multiple context mode, configure these settings in the system configuration.

Fields

- Hardware Port—*Display only.* Displays the interface ID.
- MAC Address—*Display only.* Displays the Interface MAC address.
- Media Type—Sets the media type to RJ45 or SFP. SFP is only available for SSM interfaces on the ASA 5500 series adaptive security appliance. The default is RJ45.
- Duplex—Lists the duplex options for the interface, including Full, Half, or Auto, depending on the interface type.
- Speed—Lists the speed options for the interface. The speeds available depend on the interface type. For SFP interfaces, which are always 1000 Mbps, and you can set the speed to Negotiate or Nonegotiate. Negotiate (the default) enables link negotiation, which exchanges flow-control parameters and remote fault information. Nonegotiate does not negotiate link parameters. For RJ-45 interfaces on the ASA 5500 series adaptive security appliance, the default auto-negotiation setting also includes the Auto-MDI/MDIX feature. Auto-MDI/MDIX eliminates the need for crossover cabling by performing an internal crossover when a straight cable is detected during the

auto-negotiation phase. Either the speed or duplex must be set to auto-negotiate to enable Auto-MDI/MDIX for the interface. If you explicitly set both the speed and duplex to a fixed value, thus disabling auto-negotiation for both settings, then Auto-MDI/MDIX is also disabled.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System[1]
•	•	•	•	—

1. For the system Hardware Properties dialog box, see the system Hardware Properties dialog box.

Configuring Switch Ports and VLAN Interfaces for the Cisco ASA 5505 Adaptive Security Appliance

This chapter describes how to configure the switch ports and VLAN interfaces of the ASA 5505 adaptive security appliance.

Note To configure interfaces of other models, see Chapter 4, "Configuring Interfaces."

This chapter includes the following sections:

Interface Overview

This section describes the ports and interfaces of the ASA 5505 adaptive security appliance, and includes the following topics:

Understanding ASA 5505 Ports and Interfaces

The ASA 5505 adaptive security appliance supports a built-in switch. There are two kinds of ports and interfaces that you need to configure:

- Physical switch ports—The adaptive security appliance has eight Fast Ethernet switch ports that forward traffic at Layer 2, using the switching function in hardware. Two of these ports are PoE ports. See the "For same security interfaces, you can configure established commands for both directions." section on page 5-17 for more information. You can connect these interfaces directly to user equipment such as PCs, IP phones, or a DSL modem. Or you can connect to another switch.
- Logical VLAN interfaces—In routed mode, these interfaces forward traffic between VLAN networks at Layer 3, using the configured security policy to apply firewall and VPN services. In transparent mode, these interfaces forward traffic between the VLANs on the same network at Layer 2, using the configured security policy to apply firewall services. See the "Maximum Active VLAN Interfaces for Your License" section for more information about the maximum VLAN interfaces. VLAN interfaces let you divide your equipment into separate VLANs, for example, home, business, and Internet VLANs.

To segregate the switch ports into separate VLANs, you assign each switch port to a VLAN interface. Switch ports on the same VLAN can communicate with each other using hardware switching. But when a switch port on VLAN 1 wants to communicate with a switch port on VLAN 2, then the adaptive security appliance applies the security policy to the traffic and routes or bridges between the two VLANs.

Note Subinterfaces are not available for the ASA 5505 adaptive security appliance.

Maximum Active VLAN Interfaces for Your License

In transparent firewall mode, you can configure two active VLANs in the Base license and three active VLANs in the Security Plus license, one of which must be for failover.

In routed mode, you can configure up to three active VLANs with the Base license, and up to 20 active VLANs with the Security Plus license.

An active VLAN is a VLAN with a **nameif** command configured.

With the Base license, the third VLAN can only be configured to initiate traffic to one other VLAN. See Figure 5-1 for an example network where the Home VLAN can communicate with the Internet, but cannot initiate contact with Business.

Figure 5-1 ASA 5505 Adaptive Security Appliance with Base License

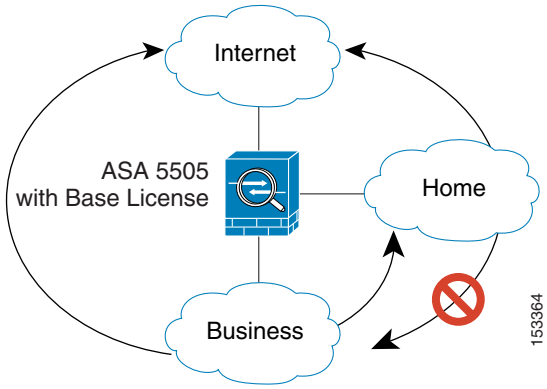

With the Security Plus license, you can configure 20 VLAN interfaces. You can configure trunk ports to accomodate multiple VLANs per port.

Note The ASA 5505 adaptive security appliance supports Active/Standby failover, but not Stateful failover.

See Figure 5-2 for an example network.

Figure 5-2 ASA 5505 Adaptive Security Appliance with Security Plus License

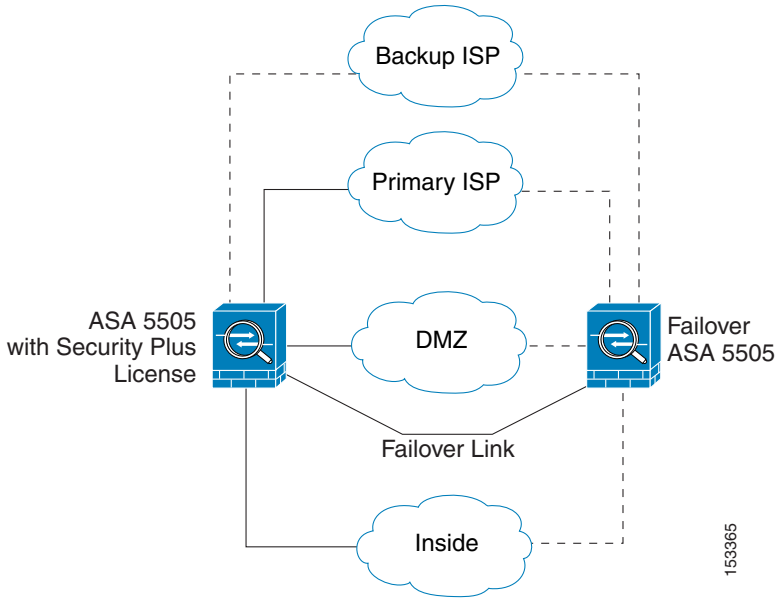

Default Interface Configuration

If your adaptive security appliance includes the default factory configuration, your interfaces are configured as follows:

- The outside interface (security level 0) is VLAN 2.

  Ethernet0/0 is assigned to VLAN 2 and is enabled.

  The VLAN 2 IP address is obtained from the DHCP server.

- The inside interface (security level 100) is VLAN 1

  Ethernet 0/1 through Ethernet 0/7 are assigned to VLAN 1 and is enabled.

  VLAN 1 has IP address 192.168.1.1.

Restore the default factory configuration using the **configure factory-default** command.

Use the procedures in this chapter to modify the default configuration, for example, to add VLAN interfaces.

If you do not have a factory default configuration, all switch ports are in VLAN 1, but no other parameters are configured.

VLAN MAC Addresses

In routed firewall mode, all VLAN interfaces share a MAC address. Ensure that any connected switches can support this scenario. If the connected switches require unique MAC addresses, you can manually assign MAC addresses.

In transparent firewall mode, each VLAN has a unique MAC address. You can override the generated MAC addresses if desired by manually assigning MAC addresses.

Power Over Ethernet

Ethernet 0/6 and Ethernet 0/7 support PoE for devices such as IP phones or wireless access points. If you install a non-PoE device or do not connect to these switch ports, the adaptive security appliance does not supply power to the switch ports.

If you shut down the switch port from the Edit Switch Port dialog box, you disable power to the device. Power is restored when you enter reenable it.

To view the status of PoE switch ports, including the type of device connected (Cisco or IEEE 802.3af), use the **show power inline** command.

Monitoring Traffic Using SPAN

If you want to monitor traffic that enters or exits one or more switch ports, you can enable SPAN, also known as switch port monitoring. The port for which you enable SPAN (called the destination port) receives a copy of every packet transmitted or received on a specified source port. The SPAN feature lets you attach a sniffer to the destination port so you can monitor all traffic; without SPAN, you would have to attach a sniffer to every port you want to monitor. You can only enable SPAN for one destination port.

See the **switchport monitor** command in the *Cisco Security Appliance Command Reference* for more information.

Security Level Overview

Each VLAN interface must have a security level in the range 0 to 100 (from lowest to highest). For example, you should assign your most secure network, such as the inside business network, to level 100. The outside network connected to the Internet can be level 0. Other networks, such as a home network can be in between. You can assign interfaces to the same security level.

The level controls the following behavior:

- Network access—By default, there is an implicit permit from a higher security interface to a lower security interface (outbound). Hosts on the higher security interface can access any host on a lower security interface. You can limit access by applying an access list to the interface.

  For same security interfaces, there is an implicit permit for interfaces to access other interfaces on the same security level or lower.

- Inspection engines—Some application inspection engines are dependent on the security level. For same security interfaces, inspection engines apply to traffic in either direction.
  - NetBIOS inspection engine—Applied only for outbound connections.
  - SQL*Net inspection engine—If a control connection for the SQL*Net (formerly OraServ) port exists between a pair of hosts, then only an inbound data connection is permitted through the adaptive security appliance.

- Filtering—HTTP(S) and FTP filtering applies only for outbound connections (from a higher level to a lower level).

  For same security interfaces, you can filter traffic in either direction.

- NAT control—When you enable NAT control, you must configure NAT for hosts on a higher security interface (inside) when they access hosts on a lower security interface (outside).

  Without NAT control, or for same security interfaces, you can choose to use NAT between any interface, or you can choose not to use NAT. Keep in mind that configuring NAT for an outside interface might require a special keyword.

- **established** command—This command allows return connections from a lower security host to a higher security host if there is already an established connection from the higher level host to the lower level host.

  For same security interfaces, you can configure **established** commands for both directions.

Configuring VLAN Interfaces

For information about how many VLANs you can configure, see the "Maximum Active VLAN Interfaces for Your License" section on page 5-14.

Note If you are using failover, do not use this procedure to name interfaces that you are reserving for failover communications. See Chapter 12, "Failover," to configure the failover link.

If you enabled Easy VPN, you cannot add or delete VLAN interfaces, nor can you edit the security level or interface name. We suggest that you finalize your interface configuration before you enable Easy VPN.

This section includes the following topics:

Interfaces > Interfaces

Configuration > Interfaces > Interfaces

The Interfaces tab displays configured VLAN interfaces. You can add or delete VLAN interfaces, and also enable communication between interfaces on the same security level or enable traffic to enter and exit the same interface.

Transparent firewall mode allows only two interfaces to pass through traffic.

Fields

- Name—Displays the interface name.
- Switch Ports—Shows the switch ports assigned to this VLAN interface.
- Enabled—Indicates if the interface is enabled, Yes or No.
- Security Level—Displays the interface security level between 0 and 100. By default, the security level is 0.
- IP Address—Displays the IP address, or in transparent mode, the word "native." Transparent mode interfaces do not use IP addresses. To set the IP address for the context or the security appliance, see the Management IP pane.
- Subnet Mask—For routed mode only. Displays the subnet mask.
- Restrict Traffic Flow—Shows if this interface is restricted from initiating contact to another VLAN.

  With the Base license, you can only configure a third VLAN if you use this option to limit it.

  For example, you have one VLAN assigned to the outside for Internet access, one VLAN assigned to an inside business network, and a third VLAN assigned to your home network. The home network does not need to access the business network, so you can use the Restrict Traffic Flow option on the home VLAN; the business network can access the home network, but the home network cannot access the business network.

  If you already have two VLAN interfaces configured with a name, be sure to enable the Restrict Traffic Flow option before you name the third interface; the adaptive security appliance does not allow three fully functioning VLAN interfaces with the Base license on the ASA 5505 adaptive security appliance.

  Note If you upgrade to the Security Plus license, you can remove this option and achieve full functionality for this interface. If you leave this option enabled, this interface continues to be limited even after upgrading.

- Backup Interface—Shows the backup ISP interface for this interface. If this interface fails, the backup interface takes over.

  The backup interface does not pass through traffic unless the default route through the primary interface fails. This option is useful for Easy VPN; when the backup interface becomes the primary, the security appliance moves the VPN rules to the new primary interface.

To ensure that traffic can pass over the backup interface in case the primary fails, be sure to configure default routes on both the primary and backup interfaces so that the backup interface can be used when the primary fails. For example, you can configure two default routes: one for the primary interface with a lower administrative distance, and one for the backup interface with a higher distance. To configure dual ISP support, see the "Static Route Tracking" section on page 14-30.

- VLAN—Shows the VLAN ID for this interface.
- Management Only—Indicates if the interface allows traffic to the security appliance or for management purposes only.
- MTU—Displays the MTU. By default, the MTU is 1500.
- Active MAC Address—Shows the active MAC address, if you assigned one manually on the Add/Edit Interface > Advanced tab.
- Standby MAC Address—Shows the standby MAC address (for failover), if you assigned one manually.
- Description—Displays a description. In the case of a failover or state link, the description is fixed as "LAN Failover Interface," "STATE Failover Interface," or "LAN/STATE Failover Interface," for example. You cannot edit this description.
- Add—Adds an interface. If you enabled Easy VPN, you cannot add VLAN interfaces.
- Edit—Edits the selected interface. If you assign an interface as the failover link or state link (see the Failover: Setup tab), you cannot edit the interface in this pane. If you enabled Easy VPN, you cannot edit the security level or interface name.
- Delete—Deletes the selected interface. If you assign an interface as the failover link or state link (see the Failover: Setup tab), you cannot delete the interface in this pane. If you enabled Easy VPN, you cannot delete VLAN interfaces.
- Enable traffic between two or more interfaces which are configured with same security levels—Enables communication between interfaces on the same security level. If you enable same security interface communication, you can still configure interfaces at different security levels as usual.
- Enable traffic between two or more hosts connected to the same interface—Enables traffic to enter and exit the same interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

Add/Edit Interface > General

Configuration > Interfaces > Add/Edit Interface > General

The Add/Edit Interface > General tab lets you add or edit a VLAN interface.

If you intend to use an interface for failover, do not configure the interface in this dialog box; instead, use the Failover: Setup tab. In particular, do not set the interface name, as this parameter disqualifies the interface from being used as the failover link; other parameters are ignored.

If you enabled Easy VPN, you cannot edit the security level or interface name. We suggest that you finalize your interface configuration before you enable Easy VPN.

After you assign the interface as the failover link or state link, you cannot edit or delete the interface from the Interfaces pane. The only exception is if you set a physical interface to be the state link, then you can configure the speed and duplex.

Fields

- Switch Ports—Assigns switch ports to this VLAN interface.
  - Available Switch Ports—Lists all switch ports, even if they are currently assigned to a different interface.
  - Selected Switch Ports—Lists the switch ports assigned to this interface.
  - Add—Adds a selected switch port to the interface. You see the following message:

    "switchport is associated with *name* interface. Adding it to this interface, will remove it from *name* interface. Do you want to continue?"

    Click **OK** to add the switch port.

    You will always see this message when adding a switch port to an interface; switch ports are assigned to the VLAN 1 interface by default even when you do not have any configuration.
  - Remove—Removes a switch port from an interface. Because the default VLAN interface for switch ports is VLAN 1, removing a switch port from an interface essentially just reassigns that switch port to VLAN 1.
- Enable Interface—Enables this interface to pass traffic. In addition to this setting, you need to set an IP address (for routed mode) and a name before traffic can pass according to your security policy.
- Dedicate this interface to management only—Sets the interface to accept traffic to the security appliance only, and not through traffic. You cannot set a primary or backup ISP interface to be management only.
- Interface Name—Sets an interface name up to 48 characters in length.
- Security Level—Sets the security level between 0 (lowest) and 100 (highest).The security appliance lets traffic flow freely from an inside network to an outside network (lower security level). Many other security features are affected by the relative security level of two interfaces.
- IP Address—For routed mode only, sets the IP address.
  - Use Static IP—Manually sets the IP address.

    IP address—Sets the IP address.

    Subnet Mask—Sets the subnet mask.
  - Obtain Address via DHCP—Dynamically sets the IP address using DHCP.

    Obtain Default Route Using DHCP—Obtains a default route from the DHCP server so that you do not need to configure a default static route.

    Renew DHCP Lease—Renews the DHCP lease.

    Retry Count—Sets the number of times between 4 and 16 that the security appliance resends a DHCP request if it does not receive a reply after the first attempt. The total number of attempts is the retry count plus the first attempt. For example, if you set the retry count to 4, the security appliance sends up to 5 DHCP requests.

DHCP Learned Route Metric—Assigns an administrative distance to the learned route. Valid values are from 1 to 255. If this field is left blank, the administrtive distance for the learned routes is 1.

Enable tracking—Check this checkbox to enable route tracking for DHCP-learned routes.

Note Route tracking is only available in single, routed mode.

Track ID—A unique identifier for the route tracking process. Valid values are from 1 to 500.

Track IP Address—Enter the IP address of the target being tracked. Typically, this would be the IP address of the next hop gateway for the route, but it could be any network object available off of that interface.

SLA ID—A unique identifier for the SLA monitoring process. Valid values are from 1 to 2147483647.

Monitor Options—Click this button to open the Route Monitoring Options dialog box. In the Route Monitoring Options dialog box you can configure the parameters of the tracked object monitoring process.

- Use PPPoE—Dynamically sets the IP address using PPPoE.

  Group Name—Specify a group name.

  PPPoE Username—Specify the username provided by your ISP.

  PPPoE Password—Specify the password provided by your ISP.

  Confirm Password—Specify the password provided by your ISP.

  PPP Authentication—Select either PAP, CHAP, or MSCHAP. PAP passes cleartext username and password during authentication and is not secure. With CHAP, the client returns the encrypted [challenge plus password], with a cleartext username in response to the server challenge. CHAP is more secure than PAP, but it does not encrypt data. MSCHAP is similar to CHAP but is more secure because the server stores and compares only encrypted passwords rather than cleartext passwords as in CHAP. MSCHAP also generates a key for data encryption by MPPE.

  Store Username and Password in Local Flash—Stores the username and password in a special location of NVRAM on the security appliance. If an Auto Update Server sends a clear config command to the security appliance and the connection is then interrupted, the security appliance can read the username and password from NVRAM and re-authenticate to the Access Concentrator.

  IP Address and Route Settings—displays the PPPoE IP Address and Route Settings dialog where you can choose addressing and tracking options. See the "PPPoE IP Address and Route Settings" section on page 4-10.

- Description—Sets an optional description up to 240 characters on a single line, without carriage returns. For multiple context mode, the system description is independent of the context description. In the case of a failover or state link, the description is fixed as "LAN Failover Interface," "STATE Failover Interface," or "LAN/STATE Failover Interface," for example. You cannot edit this description. The fixed description overwrites any description you enter here if you make this interface a failover or state link.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

Add/Edit Interface > Advanced

Configuration > Interfaces > Add/Edit Interface > Advanced

The Add/Edit Interface > Advanced tab lets you set the MTU, VLAN ID, MAC addresses, and other options.

Fields

- MTU—Sets the MTU from 300 to 65,535 bytes. The default is 1500 bytes. For multiple context mode, set the MTU in the context configuration.
- VLAN ID—Sets the VLAN ID for this interface between 1 and 1001. If you do not want to assign the VLAN ID, ASDM assigns one for you randomly.
- Mac Address Cloning—Manually assigns MAC addresses.

  By default in routed mode, all VLANs use the same MAC address. In transparent mode, the VLANs use unique MAC addresses. You might want to set unique VLANs or change the generated VLANs if your switch requires it, or for access control purposes.

  - Active Mac Address—Assigns a MAC address to the interface in H.H.H format, where H is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE.
  - Standby Mac Address—For use with failover, set the Standby Mac Address. If the active unit fails over and the standby unit becomes active, the new active unit starts using the active MAC addresses to minimize network disruption, while the old active unit uses the standby address.
- Block Traffic—Restrict this VLAN interface from initiating contact to another VLAN.

  With the Base license, you can only configure a third VLAN if you use this option to limit it.

  For example, you have one VLAN assigned to the outside for Internet access, one VLAN assigned to an inside business network, and a third VLAN assigned to your home network. The home network does not need to access the business network, so you can use the Restrict Traffic Flow option on the home VLAN; the business network can access the home network, but the home network cannot access the business network.

  If you already have two VLAN interfaces configured with a name, be sure to enable the Restrict Traffic Flow option before you name the third interface; the adaptive security appliance does not allow three fully functioning VLAN interfaces with the Base license on the ASA 5505 adaptive security appliance and will not allow you to configure one.

Note If you upgrade to the Security Plus license, you can remove this option and achieve full functionality for this interface. If you leave this option enabled, this interface continues to be limited even after upgrading.

  - Block Traffic from this Interface to—Choose a VLAN ID in the list.

- Select Backup Interface—Shows the backup ISP interface for this interface. If this interface fails, the backup interface takes over. The backup interface does not pass through traffic unless the default route through the primary interface fails. This option is useful for Easy VPN; when the backup interface becomes the primary, the security appliance moves the VPN rules to the new primary interface.

  To ensure that traffic can pass over the backup interface in case the primary fails, be sure to configure default routes on both the primary and backup interfaces so that the backup interface can be used when the primary fails. For example, you can configure two default routes: one for the primary interface with a lower administrative distance, and one for the backup interface with a higher distance. To configure dual ISP support, see the "Static Route Tracking" section on page 14-30.

  - Backup Interface—Choose a VLAN ID in the list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

Configuring Switch Ports

This section describes how to configure switch ports, and includes the following topics:

Caution The ASA 5505 adaptive security appliance does not support Spanning Tree Protocol for loop detection in the network. Therefore you must ensure that any connection with the adaptive security appliance does not end up in a network loop.

Interfaces > Switch Ports

Configuration > Interfaces > Switch Ports

The Switch Ports tab displays the switch port parameters.

Fields

- Switch Port—Lists the switch ports in the security appliance.
- Enabled—Shows if the switch port is enabled, Yes or No.
- Associated VLANs—Lists the VLAN interfaces to which the switch port is assigned. A trunk switch port can be associated with multiple VLANs.

- Associated Interface Names—Lists the VLAN interface names.
- Mode—The mode, Access or Trunk. Access ports can be assigned to one VLAN. Trunk ports can carry multiple VLANs using 802.1Q tagging. Trunk mode is available only with the Security Plus license.
- Protected—Shows if this switch port is protected, Yes or No. This option prevents the switch port from communicating with other protected switch ports on the same VLAN. You might want to prevent switch ports from communicating with each other if the devices on those switch ports are primarily accessed from other VLANs, you do not need to allow intra-VLAN access, and you want to isolate the devices from each other in case of infection or other security breach. For example, if you have a DMZ that hosts three web servers, you can isolate the web servers from each other if you apply the Protected option to each switch port. The inside and outside networks can both communicate with all three web servers, and vice versa, but the web servers cannot communicate with each other.
- Edit—Edits the switch port.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

Edit Switch Port

Configuration > Interfaces > Switch Ports > Edit Switch Port >

The Edit Switch Port dialog box lets you configure the mode, assign a switch port to a VLAN, and set the Protected option.

Fields

- Switch Port—*Display only.* Shows the selected switch port ID.
- Enable Switch Port—Enables this switch port.
- Mode and VLAN IDs—Sets the mode and the assigned VLANs.
  - Access VLAN ID—Sets the mode to access mode. Enter the VLAN ID to which you want to assign this switch port. By default, the VLAN ID is derived from the VLAN interface configuration in Interfaces > Interfaces. You can change the VLAN assignment in this dialog box. Be sure to apply the change to update the Interfaces > Interfaces tab with the new information. If you want to specify a VLAN that has not yet been added, we suggest you add the VLAN from the Interfaces > Interfaces tab and specify the switch port in the Add/Edit Interface > General tab rather than specifying it in this dialog box; in either case, you need to add the VLAN on the Interfaces > Interfaces tab and assign the switch port to it.
  - Trunk VLAN IDs—Sets the mode to trunk mode using 802.1Q tagging. Trunk mode is available only with the Security Plus license. Enter the VLAN IDs to which you want to assign this switch port, separated by commas. Trunk ports do not support untagged packets; there is no native VLAN support, and the adaptive security appliance drops all packets that do not contain a tag specified in this command. If the VLANs are already in your configuration, after you apply the

change, the Interfaces > Interfaces tab shows this switch port added to each VLAN. If you want to specify a VLAN that has not yet been added, we suggest you add the VLAN from the Interfaces > Interfaces tab and specify the switch port in the Add/Edit Interface > General tab rather than specifying it in this dialog box; in either case, you need to add the VLAN on the Interfaces > Interfaces tab and assign the switch port to it.

- Isolated—This option prevents the switch port from communicating with other protected switch ports on the same VLAN. You might want to prevent switch ports from communicating with each other if the devices on those switch ports are primarily accessed from other VLANs, you do not need to allow intra-VLAN access, and you want to isolate the devices from each other in case of infection or other security breach. For example, if you have a DMZ that hosts three web servers, you can isolate the web servers from each other if you apply the Protected option to each switch port. The inside and outside networks can both communicate with all three web servers, and vice versa, but the web servers cannot communicate with each other.
  - Isolated—Sets this switch port as a protected port.
- Duplex—Lists the duplex options for the interface, including Full, Half, or Auto. The Auto setting is the default. If you set the duplex to anything other than Auto on PoE ports Ethernet 0/6 or 0/7, then Cisco IP phones and Cisco wireless access points that do not support IEEE 802.3af will not be detected and supplied with power.
- Speed—The Auto setting is the default. If you set the speed to anything other than Auto on PoE ports Ethernet 0/6 or 0/7, then Cisco IP phones and Cisco wireless access points that do not support IEEE 802.3af will not be detected and supplied with power. The default Auto setting also includes the Auto-MDI/MDIX feature. Auto-MDI/MDIX eliminates the need for crossover cabling by performing an internal crossover when a straight cable is detected during the auto-negotiation phase. Either the speed or duplex must be set to Auto to enable Auto-MDI/MDIX for the interface. If you explicitly set both the speed and duplex to a fixed value, thus disabling auto-negotiation for both settings, then Auto-MDI/MDIX is also disabled.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

Global Objects

The Global Objects pane provides a single location where you can configure, view, and modify the reusable components that you need to implement your policy on the security appliance. For example, once you define the hosts and networks that are covered by your security policy, you can select the host or network to which a feature applies, instead of having to redefine it every time. This saves time and ensures consistency and accuracy of your security policy. When you need to add or delete a host or network, you can use the Global Objects pane to change it in a single place.

This chapter includes the following sections:

Configuring Network Object Groups

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

Network Object Groups

Configuration > Global Objects > Network Object Groups

The Network Object Groups pane lets you define network object groups. Network object groups let you predefine host and network IP addresses so that you can streamline subsequent configuration. When you configure the security policy, such as an access list or a AAA rule, you can choose these predefined addresses instead of typing them in manually. Grouping together multiple hosts and networks lets you easily apply a rule to a group of addresses.

A network object group consists of one or more IP address objects. An IP address object is either a host or subnet. You can add IP address objects manually when you create the network object group, or you can add IP address objects to the group that are derived from other configuration, such as access rules and AAA rules.

Multiple network object groups can be nested into a "group of groups" and used as a single group.

You can use a network object group for most configurations that require you to identify an IP address or network. When you are configuring NAT or security policy rules, the ASDM window even includes a side pane at the right that shows available IP address objects, network object groups, and other global objects; you can add, edit, or delete objects directly in the side pane.

Fields

- Add—Adds a network object group.
- Edit—Edits a network object group.
- Delete—Deletes a network object group. When a network object group is deleted, it is removed from all network object groups where it is used. If a network object group is used in an access rule, do not remove it. A network object group used in an access rule cannot be made empty.
- Find—Filters the display to show only matching network object groups and IP addresses. Clicking **Find** opens the Filter field. Click **Find** again to hide the Filter field.
  - Filter field—Enter the name or IP address included in the network object group. The wildcard characters asterisk (*) and question mark (?) are allowed.
  - Filter—Runs the filter.
  - Clear—Clears the Filter field.
- Name—Shows the name of the network object group. Click the plus (+) icon next to the name to expand the object group so you can view the IP addresses. Click the minus (-) icon to collapse the network object group.
- IP Address—When the network object group is expanded, shows the IP addresses.
- Netmask—When the network object group is expanded, shows the subnet masks.
- Description—Shows the description of the network object group.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Network Object Group

Configuration > Global Objects > Network Object Groups > Add/Edit Network Object Group

The Add/Edit Network Object Group lets you assign IP addresses or existing network object groups to a network object group.

Fields

- Group Name—Enter the group name, up to 64 characters in length. The name must be unique for all object groups. A network object group name cannot share a name with a service group.
- Description—Enter a description of this network object group, up to 200 characters in length.
- Members Not in Group—Lets you assign IP addresses and network object groups to the current network object group.
  - Existing Address—Lets you choose from already defined IP address objects and network object groups. Existing IP address objects are derived from other configuration, such as access rules and AAA rules.

    Name—Lists the already defined IP addresses and network object groups.
  - New Address—Lets you add a new IP address.

    IP Address—Enter an IP address.

    Netmask—Enter the subnet mask for the IP address.
- Members in Group—Shows IP addresses and network object groups that are already added to the network object group.
- Add—Adds the selected IP address to the network object group.
- Remove—Removes the selected IP address from the network object group.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Browse Address

(Various)

The Browse Address dialog box lets you choose an IP address object, IP name, or network object group. This dialog box is used in multiple configuration screens and is named appropriately for your current task. For example, from the Add/Edit Access Rule dialog box, this dialog box is named “Browse Source Address” or “Browse Destination Address.”

Fields

- Add—Adds a network object group or IP name.
- Edit—Edits the selected network object group or IP name.
- Delete—Deletes the selected network object group or IP name.
- Find—Filters the display to show only matching names or IP addresses. Clicking **Find** opens the Filter field. Click **Find** again to hide the Filter field.
  - Filter field—Enter the name or IP address included in the network object group. The wildcard characters asterisk (*) and question mark (?) are allowed.

  - Filter—Runs the filter.
  - Clear—Clears the Filter field.
- Type—Lets you choose the type of object to show, including IP Address Objects, IP Names, or Network Object Groups. To view all objects, choose **All**.
- Name—Shows the name of the object. In the case of IP address objects, the IP address is used as the name. The exception is for the "any" IP address object. Click the plus (+) icon next to the name of an item to expand it. Click the minus (-) icon to collapse the item.
- IP Address—Shows the IP addresses.
- Netmask—Shows the subnet masks.
- Description—Shows the description.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Configuring IP Names

This section describes how to associate a host IP address with a name, and includes the following topics:

- IP Names, page 6-4
- Add/Edit IP Name, page 6-5

IP Names

Configuration > Global Objects > IP Names

ASDM lets you associate a host IP address with a name, so you can use the name in your configuration instead of the IP address. You can update the IP address at any time and not disrupt your configuration. Many features on the security appliance do not support DNS, so this feature accomplishes a similar goal.

Fields

- Add—Adds an IP name.
- Edit—Edits an IP name.
- Delete—Deletes an IP name.
- Find—Filters the display to show only matching names or IP addresses. Clicking **Find** opens the Filter field. Click **Find** again to hide the Filter field.
  - Filter field—Enter the name or IP address. The wildcard characters asterisk (*) and question mark (?) are allowed.
  - Filter—Runs the filter.

  - Clear—Clears the Filter field.
- Name—Shows the names.
- IP Address—Shows the IP addresses.
- Description—Shows the description.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit IP Name

Configuration > Global Objects > IP Names > Add IP Name

The Add/Edit IP Name dialog box lets you associate an IP address with a name.

Fields

- Name—Enter the name to associate with the IP address. Use characters a to z, A to Z, 0 to 9, a dash, and an underscore. The name must be 63 characters or less. Also, the name cannot start with a number.
- IP Address—Enter the IP address.
- Description—Enter a description up to 200 characters in length.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Configuring Service Groups

This section describes how to configure service groups, and includes the following topics:

Service Groups

Configuration > Global Objects > Service Groups

The Service Groups pane lets you associate multiple services into a named group. You can create service groups for each of the following types:

- TCP ports
- UDP ports
- TCP-UDP ports
- ICMP types
- IP protocols

Multiple service groups can be nested into a "group of groups" and used as a single group.

You can use a service group for most configurations that require you to identify a port, ICMP type, or protocol. When you are configuring NAT or security policy rules, the ASDM window even includes a side pane at the right that shows available service groups and other global objects; you can add, edit, or delete objects directly in the side pane.

Fields

- Add—Adds a service group. Choose the type of service groups you want to add from the drop-down list.
- Edit—Edits a service group.
- Delete—Deletes a service group. When a service group is deleted, it is removed from all service groups where it is used. If a service group is used in an access rule, do not remove it. A service group used in an access rule cannot be made empty.
- Find—Filters the display to show only matching names. Clicking **Find** opens the Filter field. Click **Find** again to hide the Filter field.
  - Filter field—Enter the name of the service group. The wildcard characters asterisk (*) and question mark (?) are allowed.
  - Filter—Runs the filter.
  - Clear—Clears the Filter field.
- Type—Lets you choose the type of service group to show, including TCP, UDP, TCP-UDP, ICMP, and Protocol. To view all service groups, choose **All**.
- Name—Lists the service group names. Click the plus (+) icon next to the name to expand the service group so you can view the services. Click the minus (-) icon to collapse the service group.
- Description—Lists the service group descriptions.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Service Group

Configuration > Global Objects > Service Groups > Add/Edit Service Group

The Add/Edit Service Group dialog box lets you assign services to a service group. This dialog box name matches the type of service group you are adding; for example, if you are adding a TCP service group, the name is "Add/Edit TCP Service Group."

Fields

- Group Name—Enter the group name, up to 64 characters in length. The name must be unique for all object groups. A service group name cannot share a name with a network object group.
- Description—Enter a description of this service group, up to 200 characters in length.
- Members Not in Group—Identifies items that can be added to the service group.
  - Service/Service Group, ICMP Type/ICMP Group, or Protocol/Protocol Group—The title of this table depends on the type of service group you are adding. Choose from already defined service groups, or choose from a list of commonly used port, type, or protocol names.

    Name—Lists the already defined service groups and commonly used ports, types, or protocols.
  - Port #, ICMP #, or Protocol #—The title of this table depends on the type of service group you are adding. Lets you add a new item, either by number or name. For TCP, UDP, and TCP-UDP service groups, you can enter a range of ports numbers.
- Members in Group—Shows items that are already added to the service group.
- Add—Adds the selected item to the service group.
- Remove—Removes the selected item from the service group.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Browse Service Groups

(Various)

The Browse Service Groups dialog box lets you choose a service group. This dialog box is used in multiple configuration screens and is named appropriately for your current task. For example, from the Add/Edit Access Rule dialog box, this dialog box is named "Browse Source Port" or "Browse Destination Port."

Fields

- Add—Adds a service group.
- Edit—Edits the selected service group.
- Delete—Deletes the selected service group.

- Find—Filters the display to show only matching names. Clicking **Find** opens the Filter field. Click **Find** again to hide the Filter field.
  - Filter field—Enter the name of the service group. The wildcard characters asterisk (*) and question mark (?) are allowed.
  - Filter—Runs the filter.
  - Clear—Clears the Filter field.
- Type—Lets you choose the type of service group to show, including TCP, UDP, TCP-UDP, ICMP, and Protocol. To view all types, choose **All**. Typically, the type of rule you configure can only use one type of service group; you cannot select a UDP service group for a TCP access rule.
- Name—Shows the name of the service group. Click the plus (+) icon next to the name of an item to expand it. Click the minus (-) icon to collapse the item.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Configuring Class Maps

An inspection class map matches application traffic with criteria specific to the application, such as a URL string. 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.

This section describes how to configure inspection class maps, and includes the following topics:

DNS Class Map

Configuration > Global Objects > Class Maps > DNS

The DNS Class Map panel lets you configure DNS class maps for DNS 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 DNS 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 DNS class map.
  - Value—Shows the value to match in the DNS class map.
- Description—Shows the description of the class map.
- Add—Adds match conditions for the DNS class map.
- Edit—Edits match conditions for the DNS class map.
- Delete—Deletes match conditions for the DNS class map.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit DNS Traffic Class Map

Configuration > Global Objects > Class Maps > DNS > Add/Edit DNS Traffic Class Map

The Add/Edit DNS Traffic Class Map dialog box lets you define a DNS class map.

Fields

- Name—Enter the name of the DNS class map, up to 40 characters in length.
- Description—Enter the description of the DNS class map.
- Add—Adds a DNS class map.
- Edit—Edits a DNS class map.
- Delete—Deletes a DNS class map.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit DNS Match Criterion

Configuration > Global Objects > Class Maps > DNS > Add/Edit DNS Traffic Class Map > Add/Edit DNS Match Criterion

The Add/Edit DNS Match Criterion dialog box lets you define the match criterion and value for the DNS 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 DNS traffic to match.
  - Header Flag—Match a DNS flag in the header.
  - Type—Match a DNS query or resource record type.
  - Class—Match a DNS query or resource record class.
  - Question—Match a DNS question.
  - Resource Record—Match a DNS resource record.
  - Domain Name—Match a domain name from a DNS query or resource record.
- Header Flag Criterion Values—Specifies the value details for the DNS header flag match.
  - Match Option—Specifies either an exact match or match all bits (bit mask match).
  - Match Value—Specifies to match either the header flag name or the header flag value.

    Header Flag Name—Lets you select one or more header flag names to match, including AA (authoritative answer), QR (query), RA (recursion available), RD (recursion denied), TC (truncation) flag bits.

    Header Flag Value—Lets you enter an arbitrary 16-bit value in hex to match.
- Type Criterion Values—Specifies the value details for the DNS type match.
  - DNS Type Field Name—Lists the DNS types to select.

    A—IPv4 address

    NS—Authoritative name server

    CNAME—Canonical name

    SOA—Start of a zone of authority

    TSIG—Transaction signature

    IXFR—Incremental (zone) transfer

AXFR—Full (zone) transfer

- DNS Type Field Value—Specifies to match either a DNS type field value or a DNS type field range.

  Value—Lets you enter an arbitrary value between 0 and 65535 to match.

  Range—Lets you enter a range match. Both values between 0 and 65535.

- Class Criterion Values—Specifies the value details for the DNS class match.
  - DNS Class Field Name—Specifies to match on internet, the DNS class field name.
  - DNS Class Field Value—Specifies to match either a DNS class field value or a DNS class field range.

    Value—Lets you enter an arbitrary value between 0 and 65535 to match.

    Range—Lets you enter a range match. Both values between 0 and 65535.
- Question Criterion Values—Specifies to match on the DNS question section.
- Resource Record Criterion Values—Specifies to match on the DNS resource record section.
  - Resource Record— Lists the sections to match.

    Additional—DNS additional resource record

    Answer—DNS answer resource record

    Authority—DNS authority resource record
- Domain Name Criterion Values—Specifies to match on the DNS domain name.
  - 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Manage Regular Expressions

Configuration > Global Objects > Class Maps > DNS > Add/Edit DNS Traffic Class Map > Add/Edit DNS Match Criterion > Manage Regular Expressions

The Manage Regular Expressions dialog box lets you configure Regular Expressions for use in pattern matching. Regular expressions that start with "_default" are default regular expressions and cannot be modified or deleted.

Fields

- Name—Shows the regular expression names.
- Value—Shows the regular expression definitions.
- Add—Adds a regular expression.
- Edit—Edits a regular expression.
- Delete—Deletes a regular expression.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Manage Regular Expression Class Maps

Configuration > Global Objects > Class Maps > DNS > Add/Edit DNS Traffic Class Map > Add/Edit DNS Match Criterion > Manage Regular Expression Class Maps

The Manage Regular Expression Class Maps dialog box lets you configure regular expression class maps. See Regular Expressions for more information.

Fields

- Name—Shows the regular expression class map name.
- Match Conditions—Shows the match type and regular expressions in the class map.
  - Match Type—Shows the match type, which for regular expressions is always a positive match type (shown by the icon with the equal sign (=)) the criteria. (Inspection class maps allow you to create negative matches as well (shown by the icon with the red circle)). If more than one regular expression is in the class map, then each match type icon appears with “OR” next it, to indicate that this class map is a “match any” class map; traffic matches the class map if only one regular expression is matched.
  - Regular Expression—Lists the regular expressions included in each class map.
- Description—Shows the description of the class map.
- Add—Adds a regular expression class map.
- Edit—Edits a regular expression class map.
- Delete—Deletes a regular expression class map.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

FTP Class Map

Configuration > Global Objects > Class Maps > FTP

The FTP Class Map panel 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit FTP Traffic Class Map

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit FTP Match Criterion

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

H.323 Class Map

Configuration > Global Objects > Class Maps > H.323

The H.323 Class Map panel 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

HTTP Class Map

Configuration > Global Objects > Class Maps > HTTP

The HTTP Class Map panel 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit HTTP Traffic Class Map

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit HTTP Match Criterion

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

IM Class Map

Configuration > Global Objects > Class Maps > Instant Messaging (IM)

The IM Class Map panel lets you configure IM class maps for IM 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 IM 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 IM class map.
  - Value—Shows the value to match in the IM class map.
- Description—Shows the description of the class map.
- Add—Adds an IM class map.
- Edit—Edits an IM class map.
- Delete—Deletes an IM class map.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit IM Traffic Class Map

Configuration > Global Objects > Class Maps > IM > Add/Edit IM Traffic Class Map

The Add/Edit IM Traffic Class Map dialog box lets you define a IM class map.

Fields

- Name—Enter the name of the IM class map, up to 40 characters in length.
- Description—Enter the description of the IM class map.
- Add—Adds an IM class map.
- Edit—Edits an IM class map.
- Delete—Deletes an IM class map.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit IM Match Criterion

Configuration > Global Objects > Class Maps > IM > Add/Edit IM Traffic Class Map > Add/Edit IM Match Criterion

The Add/Edit IM Match Criterion dialog box lets you define the match criterion and value for the IM 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 IM traffic to match.
  - Protocol—Match IM protocols.
  - Service—Match IM services.
  - Version—Match IM file transfer service version.
  - Client Login Name—Match client login name from IM service.
  - Client Peer Login Name—Match client peer login name from IM service.
  - Source IP Address—Match source IP address.
  - Destination IP Address—Match destination IP address.
  - Filename—Match filename form IM file transfer service.
- Protocol Criterion Values—Specifies which IM protocols to match.
  - Yahoo! Messenger—Specifies to match Yahoo! Messenger instant messages.
  - MSN Messenger—Specifies to match MSN Messenger instant messages.
- Service Criterion Values—Specifies which IM services to match.
  - Chat—Specifies to match IM message chat service.
  - Conference—Specifies to match IM conference service.
  - File Transfer—Specifies to match IM file transfer service.
  - Games—Specifies to match IM gaming service.
  - Voice Chat—Specifies to match IM voice chat service (not available for Yahoo IM)
  - Web Cam—Specifies to match IM webcam service.
- Version Criterion Values—Specifies to match the version from the IM file transfer service. 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.
- Client Login Name Criterion Values—Specifies to match the client login name from the IM service. 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.

- Client Peer Login Name Criterion Values—Specifies to match the client peer login name from the IM service. 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.
- Source IP Address Criterion Values—Specifies to match the source IP address of the IM service.
  - IP Address—Enter the source IP address of the IM service.
  - IP Mask—Mask of the source IP address.
- Destination IP Address Criterion Values—Specifies to match the destination IP address of the IM service.
  - IP Address—Enter the destination IP address of the IM service.
  - IP Mask—Mask of the destination IP address.
- Filename Criterion Values—Specifies to match the filename from the IM file transfer service. 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

SIP Class Map

Configuration > Global Objects > Class Maps > SIP

The SIP Class Map panel 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Configuring Inspect Maps

This section describes how to configure inspect maps, and includes the following topics:

The algorithm the security appliance uses for stateful application inspection ensures the security of applications and services. Some applications require special handling, and specific application inspection engines are provided for this purpose. Applications that require special application inspection engines are those that embed IP addressing information in the user data packet or open secondary channels on dynamically assigned ports.

Application inspection engines work with NAT to help identify the location of embedded addressing information. This allows NAT to translate these embedded addresses and to update any checksum or other fields that are affected by the translation.

Each application inspection engine also monitors sessions to determine the port numbers for secondary channels. Many protocols open secondary TCP or UDP ports to improve performance. The initial session on a well-known port is used to negotiate dynamically assigned port numbers. The application inspection engine monitors these sessions, identifies the dynamic port assignments, and permits data exchange on these ports for the duration of the specific session.

In addition, stateful application inspection audits the validity of the commands and responses within the protocol being inspected. The security appliance helps to prevent attacks by verifying that traffic conforms to the RFC specifications for each protocol that is inspected.

The Inspect Maps feature lets you create inspect maps for specific protocol inspection engines. You use an inspect map to store the configuration for a protocol inspection engine. You then enable the configuration settings in the inspect map by associating the map with a specific type of traffic using a global security policy or a security policy for a specific interface.

Use the Service Policy Rules tab on the Security Policy pane to apply the inspect map to traffic matching the criteria specified in the service policy. A service policy can apply to a specific interface or to all the interfaces on the security appliance.

DCERPC	The DCERPC inspection lets you create, view, and manage DCERPC inspect maps. You can use a DCERPC map to inspect DCERPC messages between a client and endpoint mapper, and to apply NAT for the secondary connection, if needed. DCERPC is a specification for a remote procedure call mechanism.
DNS	The DNS inspection lets you create, view, and manage DNS inspect maps. You can use a DNS map to have more control over DNS messages and to protect against DNS spoofing and cache poisoning. DNS is used to resolve information about domain names, including IP addresses and mail servers.
ESMTP	The ESMTP inspection lets you create, view, and manage ESMTP inspect maps. You can use an ESMTP map for application security and protocol conformance to protect against attacks, to block senders and receivers, and to block mail relay. Extended SMTP defines protocol extensions to the SMTP standard.
FTP	The FTP inspection lets you create, view, and manage FTP inspect maps. FTP is a common protocol used for transferring files over a TCP/IP network, such as the Internet. You can use an FTP map to block specific FTP protocol methods, such as an FTP PUT, from passing through the security appliance and reaching your FTP server.
GTP	The GTP inspection lets you create, view, and manage GTP inspect maps. 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.
H.323	The H.323 inspection lets you create, view, and manage H.323 inspect maps. You can use an H.323 map to inspect RAS, H.225, and H.245 VoIP protocols, and for state tracking and filtering.
HTTP	The HTTP inspection lets you create, view, and manage HTTP inspect maps. HTTP is the protocol used for communication between Worldwide Web clients and servers. You can use an HTTP map to enforce RFC compliance and HTTP payload content type. You can also block specific HTTP methods and prevent the use of certain tunneled applications that use HTTP as the transport.
IM	The IM inspection lets you create, view, and manage IM inspect maps. You can use an IM map to control the network usage and stop leakage of confidential data and other network threats from IM applications.
IPSec Pass Through	The IPSec Pass Through inspection lets you create, view, and manage IPSec Pass Through inspect maps. You can use an IPSec Pass Through map to permit certain flows without using an access list.
MGCP	The MGCP inspection lets you create, view, and manage MGCP inspect maps. You can use an MGCP map to manage connections between VoIP devices and MGCP call agents.

NetBIOS	The NetBIOS inspection lets you create, view, and manage NetBIOS inspect maps. You can use a NetBIOS map to enforce NetBIOS protocol conformance including field count and length consistency, and message checks.
RADIUS Accounting	The RADIUS Accounting inspection lets you create, view, and manage RADIUS Accounting inspect maps. You can use a RADIUS map to protect against an overbilling attack.
SCCP (Skinny)	The SCCP (Skinny) inspection lets you create, view, and manage SCCP (Skinny) inspect maps. You can use an SCCP map to perform protocol conformance checks and basic state tracking.
SIP	The SIP inspection lets you create, view, and manage SIP inspect maps. You can use a SIP map for application security and protocol conformance to protect against SIP-based attacks. SIP is a protocol widely used for internet conferencing, telephony, presence, events notification, and instant messaging.
SNMP	The SNMP inspection lets you create, view, and manage SNMP inspect maps. SNMP is a protocol used for communication between network management devices and network management stations. You can use an SNMP map to block a specific SNMP version, including SNMP v1, 2, 2c and 3.

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. 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.

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.
- 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 Customize Security Level dialog box for additional settings.
- Default Level—Sets the security level back to the default level of Medium.

- DCERPC Inspect Maps—Table that lists the defined DCERPC inspect maps. The defined inspect maps are also listed in the DCERPC area of the Inspect Maps tree.
- Add—Adds the new DCERPC inspect map to the defined list in the DCERPC Inspect Maps table and to the DCERPC area of the Inspect Maps tree. To configure the new DCERPC map, select the DCERPC entry in Inspect Maps tree.
- Delete—Deletes the application inspection map selected in the DCERPC Inspect Maps table and from the DCERPC area of the Inspect Maps tree.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Customize Security Level

Configuration > Global Objects > Inspect Maps > DCERPC > Customize Security Level

The Customize Security Level dialog box lets you configure the security settings for previously configured DCERPC application inspection maps.

Fields

- Settings—Specifies the pinhole timeout and endpoint mapper security settings.
  - Pinhole Timeout—Sets the pinhole timeout. Since 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.

  - Enforce endpoint-mapper service lookup—Enables the lookup operation of the endpoint mapper service. If disabled, the pinhole timeout is used.

    Service Lookup Timeout—Sets the timeout for pinholes from lookup operation.

- Reset to Predefined Security Level—Resets the security level settings to the predefined levels of high, medium, or low.
  - Reset To—Resets the security level to high, medium, or low.
- Reset—Resets all security settings to the default. The default pinhole timeout is one minute. The default endpoint mapper settings are none.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

DCERPC Inspect Map Basic/Advanced Viewl

Configuration > Global Objects > Inspect Maps > DCERPC > DCERPC Inspect Map > Basic/Advanced View

The DCERPC map pane lets you configure basic and advanced settings for previously configured DCERPC application inspection maps.

Fields

- Name—Shows the name of the previously configured DCERPC map.
- Description—Enter the description of the DCERPC map, up to 200 characters in length.
- Basic View—Shows the current security settings.
  - Customize—Opens the Customize Security Level dialog box to configure security settings.
  - Default Level—Sets the security level back to the default level of Medium.
- Advanced View—Lets you configure the security settings.
  - Pinhole Timeout—Sets the pinhole timeout. Since 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.
  - Enforce endpoint-mapper service lookup—Enables the lookup operation of the endpoint mapper service. If disabled, the pinhole timeout is used.

    Service Lookup Timeout—Sets the timeout for pinholes from lookup operation.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

DNS Inspect Map

Configuration > Global Objects > Inspect Maps > DNS

The DNS pane lets you view previously configured DNS application inspection maps. A DNS map lets you change the default configuration values used for DNS application inspection.

DNS application inspection supports DNS message controls that provide protection against DNS spoofing and cache poisoning. User configurable rules allow certain DNS types to be allowed, dropped, and/or logged, while others are blocked. Zone transfer can be restricted between servers with this function, for example.

The Recursion Desired and Recursion Available flags in the DNS header can be masked to protect a public server from attack if that server only supports a particular internal zone. In addition, DNS randomization can be enabled avoid spoofing and cache poisoning of servers that either do not support randomization, or utilize a weak pseudo random number generator. Limiting the domain names that can be queried also restricts the domain names which can be queried, which protects the public server further.

A configurable DNS mismatch alert can be used as notification if an excessive number of mismatching DNS responses are received, which could indicate a cache poisoning attack. In addition, a configurable check to enforce a Transaction Signature be attached to all DNS messages is also supported.

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.
- Security Level—Select the security level (high, medium, or low).
  - Low—Default.

    DNS Guard: enabled

    NAT rewrite: enabled

    Protocol enforcement: enabled

    ID randomization: disabled

    Message length check: enabled

    Message length maximum: 512

    Mismatch rate logging: disabled

    TSIG resource record: not enforced
  - Medium

    DNS Guard: enabled

    NAT rewrite: enabled

    Protocol enforcement: enabled

ID randomization: enabled

Message length check: enabled

Message length maximum: 512

Mismatch rate logging: enabled

TSIG resource record: not enforced

– High

DNS Guard: enabled

NAT rewrite: enabled

Protocol enforcement: enabled

ID randomization: enabled

Message length check: enabled

Message length maximum: 512

Mismatch rate logging: enabled

TSIG resource record: enforced

– Customize—Opens the Customize Security Level dialog box for additional settings.

– Default Level—Sets the security level back to the default level of Low.

- DNS Inspect Maps—Table that lists the defined DNS inspect maps. The defined inspect maps are also listed in the DNS area of the Inspect Maps tree.
- Add—Adds the new DNS inspect map to the defined list in the DNS Inspect Maps table and to the DNS area of the Inspect Maps tree. To configure the new DNS map, select the DNS entry in Inspect Maps tree.
- Delete—Deletes the application inspection map selected in the DNS Inspect Maps table and from the DNS area of the Inspect Maps tree.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Customize Security Level

Configuration > Global Objects > Inspect Maps > DNS > Customize Security Level

The Customize Security Level dialog box lets you configure the security settings for previously configured DNS application inspection maps.

Fields

- Settings—Specifies DNS security settings and actions.

  - Enable DNS guard function—As part of protocol conformance, this option performs a DNS query and response mismatch check using the identification field in the DNS header. One response per query is allowed to go through the security appliance.
  - Enable NAT rewrite function—As part of protocol conformance, this option enables IP address translation in the A record of the DNS response.
  - Enable protocol enforcement—As part of protocol conformance, this option enables DNS message format check, including domain name, label length, compression, and looped pointer check.
  - Randomize the DNS identifier for DNS query—As part of protocol conformance, this option randomizes the DNS identifier in the DNS query message.
  - Drop packets that exceed specified maximum length—As part of filtering, this option drops packets that exceed maximum length in bytes.

    Maximum Packet Length—Enter maximum packet length in bytes.
  - Enable Logging when DNS ID mismatch rate exceeds specified rate—Reports excessive instances of DNS identifier mismatches.

    Mismatch Instance Threshold—Enter the maximum number of mismatch instances before a system message log is sent.

    Time Interval—Enter the time period to monitor (in seconds).
  - Enforce TSIG record source to be present in DNS message—As part of protocol conformance, this option requires that a TSIG resource record be present in DNS transactions. Actions taken when TSIG is enforced:

    Drop packet—Drops the packet (logging can be either enabled or disabled).

    Log—Enables logging.

- Reset to predefined security level—Resets the security level settings to the predefined levels of high, medium, or low. Default is low.
  - Reset to—Specifies high, medium, or low security setting.
  - Reset—Reset settings to selected level.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

DNS Inspect Map Basic View

Configuration > Global Objects > Inspect Maps > DNS > DNS Inspect Map > Basic View

The DNS Inspect Map Basic View pane shows the configured settings for the DNS inspect map. The Advanced View lets you configure the settings.

Fields

- Name—Shows the name of the previously configured DNS map.
- Description—Enter the description of the DNS map, up to 200 characters in length.
- Security Level—Shows the current security settings.
  - Customize—Opens the Customize Security Level dialog box to configure the security settings.
  - Default Level—Sets the security level back to the default.
- Advanced View—Lets you configure the security settings.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

DNS Inspect Map Advanced View

Configuration > Global Objects > Inspect Maps > DNS > DNS Inspect Map > Advanced View

The DNS Inspect Map Advanced View pane lets you configure the inspect map settings.

Fields

- Name—Shows the name of the previously configured DNS map.
- Description—Enter the description of the DNS map, up to 200 characters in length.
- Protocol Conformance—Tab that lets you configure the protocol conformance settings for DNS.
  - Enable DNS guard function—Performs a DNS query and response mismatch check using the identification field in the DNS header. One response per query is allowed to go through the security appliance.
  - Enable NAT re-write function—Enables IP address translation in the A record of the DNS response.
  - Enable protocol enforcement—Enables DNS message format check, including domain name, label length, compression, and looped pointer check.
  - Randomize the DNS identifier for DNS query— Randomizes the DNS identifier in the DNS query message.
  - Enforce TSIG resource record to be present in DNS message—Requires that a TSIG resource record be present in DNS transactions. Actions taken when TSIG is enforced:

    Drop packet—Drops the packet (logging can be either enabled or disabled).

    Log—Enables logging.
- Filtering—Tab that lets you configure the filtering settings for DNS.
  - Global Settings—Applies settings globally.

Drop packets that exceed specified maximum length (global)—Drops packets that exceed maximum length in bytes.

Maximum Packet Length—Enter maximum packet length in bytes.

– Server Settings—Applies settings on the server only.

Drop packets that exceed specified maximum length———Drops packets that exceed maximum length in bytes.

Maximum Packet Length—Enter maximum packet length in bytes.

Drop packets sent to server that exceed length indicated by the RR—Drops packets sent to the server that exceed the length indicated by the Resource Record.

– Client Settings—Applies settings on the client only.

Drop packets that exceed specified maximum length———Drops packets that exceed maximum length in bytes.

Maximum Packet Length—Enter maximum packet length in bytes.

Drop packets sent to client that exceed length indicated by the RR—Drops packets sent to the client that exceed the length indicated by the Resource Record.

- Mismatch Rate—Tab that lets you configure the ID mismatch rate for DNS.
  - Enable Logging when DNS ID mismatch rate exceeds specified rate—Reports excessive instances of DNS identifier mismatches.

    Mismatch Instance Threshold—Enter the maximum number of mismatch instances before a system message log is sent.

    Time Interval—Enter the time period to monitor (in seconds).
- Inspections—Tab that shows you the DNS 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 DNS inspection.
  - Value—Shows the value to match in the DNS inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add DNS Inspect dialog box to add a DNS inspection.
  - Edit—Opens the Edit DNS Inspect dialog box to edit a DNS inspection.
  - Delete—Deletes a DNS inspection.
  - Move Up—Moves an inspection up in the list.
  - Move Down—Moves an inspection down in the list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit DNS Inspect

Configuration > Global Objects > Inspect Maps > DNS > DNS Inspect Map > Advanced View > Add/Edit DNS Inspect

The Add/Edit DNS Inspect dialog box lets you define the match criterion and value for the DNS inspect map.

Fields

- Single Match—Specifies that the DNS 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 DNS traffic to match.
  - Header Flag—Match a DNS flag in the header.
  - Type—Match a DNS query or resource record type.
  - Class—Match a DNS query or resource record class.
  - Question—Match a DNS question.
  - Resource Record—Match a DNS resource record.
  - Domain Name—Match a domain name from a DNS query or resource record.
- Header Flag Criterion Values—Specifies the value details for DNS header flag match.
  - Match Option—Specifies either an exact match or match all bits (bit mask match).
  - Match Value—Specifies to match either the header flag name or the header flag value.

    Header Flag Name—Lets you select one or more header flag names to match, including AA (authoritative answer), QR (query), RA (recursion available), RD (recursion denied), TC (truncation) flag bits.

    Header Flag Value—Lets you enter an arbitrary 16-bit value in hex to match.
- Type Criterion Values—Specifies the value details for DNS type match.
  - DNS Type Field Name—Lists the DNS types to select.

    A—IPv4 address

    NS—Authoritative name server

    CNAME—Canonical name

    SOA—Start of a zone of authority

    TSIG—Transaction signature

    IXFR—Incremental (zone) transfer

    AXFR—Full (zone) transfer
  - DNS Type Field Value—Specifies to match either a DNS type field value or a DNS type field range.

    Value—Lets you enter an arbitrary value between 0 and 65535 to match.

    Range—Lets you enter a range match. Both values between 0 and 65535.
- Class Criterion Values—Specifies the value details for DNS class match.

  - DNS Class Field Name—Specifies to match on internet, the DNS class field name.
  - DNS Class Field Value—Specifies to match either a DNS class field value or a DNS class field range.

    Value—Lets you enter an arbitrary value between 0 and 65535 to match.

    Range—Lets you enter a range match. Both values between 0 and 65535.
- Question Criterion Values—Specifies to match on the DNS question section.
- Resource Record Criterion Values—Specifies to match on the DNS resource record section.
  - Resource Record— Lists the sections to match.

    Additional—DNS additional resource record

    Answer—DNS answer resource record

    Authority—DNS authority resource record
- Domain Name Criterion Values—Specifies to match on DNS domain name.
  - 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 DNS inspection.
  - DNS Traffic Class—Specifies the DNS traffic class match.
  - Manage—Opens the Manage DNS Class Maps dialog box to add, edit, or delete DNS Class Maps.
- Actions—Primary action and log settings.
  - Primary Action—Mask, drop packet, drop connection, none.
  - Log—Enable or disable.
  - Enforce TSIG—Do not enforce, drop packet, log, drop packet and log.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Manage Class Maps

Configuration > Global Objects > Inspect Maps > DNS > DNS Inspect Map > Advanced View > Add DNS Inspect > Multiple Matches > Manage DNS Class Maps

The Manage Class Map dialog box lets you configure class maps for 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, Instant Messaging (IM), and SIP.

Fields

- Name—Shows the 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 class map.
  - Value—Shows the value to match in the class map.
- Description—Shows the description of the class map.
- Add—Adds match conditions for the class map.
- Edit—Edits match conditions for the class map.
- Delete—Deletes match conditions for the class map.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

ESMTP Inspect Map

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

- 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.
- 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

- Customize—Opens the Customize Security Level dialog box for additional settings.
- Default Level—Sets the security level back to the default level of Low.
- MIME File Type Filtering—Opens the MIME Type Filtering dialog box to configure MIME file type filters.

- ESMTP Inspect Maps—Table that lists the defined ESMTP inspect maps. The defined inspect maps are also listed in the ESMTP area of the Inspect Maps tree.
- Add—Adds the new ESMTP inspect map to the defined list in the ESMTP Inspect Maps table and to the ESMTP area of the Inspect Maps tree. To configure the new ESMTP map, select the ESMTP entry in Inspect Maps tree.
- Delete—Deletes the application inspection map selected in the ESMTP Inspect Maps table and from the ESMTP area of the Inspect Maps tree.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Customize Security Level

Configuration > Global Objects > Inspect Maps > ESMTP > Customize Security Level

The Customize Security Level dialog box lets you configure the security settings for previously configured ESMTP application inspection maps.

Fields

- Settings—Specifies ESMTP security settings and actions.
  - 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.
  - Check for command line length—Enables command line length matching at specified length.

    Minimum Length—Shows the minimum length configured.

    Action—Reset, drop connection, log.

    Log—Enable or disable.
  - Check for command recipient count—Enables command recipient count matching at specified count.

    Minimum Count—Shows the minimum count configured.

    Action—Reset, drop connection, log.

    Log—Enable or disable.
  - Check for body line length—Enables body line length matching at specified length.

    Minimum Length—Shows the minimum length configured.

    Action—Reset, drop connection, log.

    Log—Enable or disable.
  - Check for sender address length—Enables sender address length matching at specified length.

    Minimum Length—Shows the minimum length configured.

    Action—Reset, drop connection, log.

    Log—Enable or disable.
  - Check for MIME file name length—Enables MIME file name length matching at specified length.

    Minimum Length—Shows the minimum length configured.

    Action—Reset, drop connection, log.

    Log—Enable or disable.
- Reset to predefined security level—Resets the security level settings to the predefined levels of high, medium, or low. Default is low.
  - Reset to—Specifies high, medium, or low security setting.
  - Reset—Reset settings to selected level.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

MIME File Type Filtering

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

ESMTP Inspect Map Basic View

Configuration > Global Objects > Inspect Maps > ESMTP > ESMTP Inspect Map > Basic View

The ESMTP Inspect Map Basic View pane shows the configured settings for the ESMTP inspect map. The Advanced View lets you configure the settings.

Fields

- Name—Shows the name of the previously configured ESMTP map.
- Description—Enter the description of the ESMTP map, up to 200 characters in length.
- Security Level—Shows the current security settings.

  - Customize—Opens the Customize Security Level dialog box to configure the security settings.
  - Default Level—Sets the security level back to the default.
- MIME File Type Filtering—Opens the MIME Type Filtering dialog box to configure MIME file type filters.
- Advanced View—Lets you configure the security settings.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

ESMTP Inspect Map Advanced View

Configuration > Global Objects > Inspect Maps > ESMTP > ESMTP Inspect Map > Advanced View

The ESMTP Inspect Map Advanced View pane lets you configure the settings for the inspect map.

Fields

- Name—Shows the name of the previously configured ESMTP map.
- Description—Enter the description of the ESMTP map, up to 200 characters in length.
- 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit ESMTP Inspect

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

FTP Inspect Map

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

- 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.
- Security Level—Select the security level (medium or low).
  - Low

    Mask Banner Disabled

    Mask Reply Disabled
  - Medium—Default.

    Mask Banner Enabled

    Mask Reply Enabled
  - Customize—Opens the Customize Security Level dialog box for additional settings.
  - Default Level—Sets the security level back to the default level of Medium.
  - File Type Filtering—Opens the Type Filtering dialog box to configure file type filters.
- FTP Inspect Maps—Table that lists the defined FTP inspect maps. The defined inspect maps are also listed in the FTP area of the Inspect Maps tree.
- Add—Adds the new FTP inspect map to the defined list in the FTP Inspect Maps table and to the FTP area of the Inspect Maps tree. To configure the new FTP map, select the FTP entry in Inspect Maps tree.
- Delete—Deletes the application inspection map selected in the FTP Inspect Maps table and from the FTP area of the Inspect Maps tree.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Customize Security Level

Configuration > Global Objects > Inspect Maps > FTP > Customize Security Level

The Customize Security Level dialog box lets you configure the security settings for previously configured FTP application inspection maps.

Fields

- Settings—Specifies FTP security settings and actions.
  - 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.
- Reset to predefined security level—Resets the security level settings to the predefined levels of high, medium, or low. Default is medium.
  - Reset to—Specifies high, medium, or low security setting.
  - Reset—Reset settings to selected level.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

File Type Filtering

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

FTP Inspect Map Basic View

Configuration > Global Objects > Inspect Maps > FTP > FTP Inspect Map > Basic View

The FTP Inspect Map Basic View pane shows the configured settings for the FTP inspect map. The Advanced View lets you configure the settings.

Fields

- Name—Shows the name of the previously configured FTP map.
- Description—Enter the description of the FTP map, up to 200 characters in length.
- Security Level—Shows the current security settings.
  - Customize—Opens the Customize Security Level dialog box to configure the security settings.
  - Default Level—Sets the security level back to the default.
- File Type Filtering—Opens the Type Filtering dialog box to configure file type filters.
- Advanced View—Lets you configure the security settings.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

FTP Inspect Map Advanced View

Configuration > Global Objects > Inspect Maps > FTP > FTP Inspect Map > Advanced View

The FTP Inspect Map Advanced View pane lets you configure the settings for the inspect map.

Fields

- Name—Shows the name of the previously configured FTP map.
- Description—Enter the description of the FTP map, up to 200 characters in length.
- 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit FTP Map

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 DNS 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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

- 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.
- 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.
- Customize—Opens the Customize Security Level dialog box for additional settings.
- Default Level—Sets the security level back to the default.
- IMSI Prefix Filtering—Opens the IMSI Prefix Filtering dialog box to configure IMSI prefix filters.
- GTP Inspect Maps—Table that lists the defined GTP inspect maps. The defined inspect maps are also listed in the GTP area of the Inspect Maps tree.
- Add—Adds the new GTP inspect map to the defined list in the GTP Inspect Maps table and to the GTP area of the Inspect Maps tree. To configure the new GTP map, select the GTP entry in Inspect Maps tree.
- Delete—Deletes the application inspection map selected in the GTP Inspect Maps table and from the GTP area of the Inspect Maps tree.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Customize Security Level

Configuration > Global Objects > Inspect Maps > FTP > Customize Security Level

The Customize Security Level dialog box lets you configure the security settings for previously configured FTP application inspection maps.

Fields

- Permit Errors—Lets any packets that are invalid or that encountered an error during inspection to be sent through the security appliance instead of being dropped. By default, all invalid packets or packets that failed during parsing are dropped.
- Drop and Log unknown message IDs—Drops and logs all message IDs that are unknown.
- 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.

- Reset to—Specifies low security setting.
- Reset—Reset settings to selected level.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

GTP Inspect Map Basic View

Configuration > Global Objects > Inspect Maps > GTP > GTP Inspect Map > Basic View

The GTP Inspect Map Basic View pane shows the configured settings for the GTP inspect map. The Advanced View lets you configure the settings.

Fields

- Name—Shows the name of the previously configured GTP map.
- Description—Enter the description of the GTP map, up to 200 characters in length.
- Security Level—Shows the current security settings.
  - Customize—Opens the Customize Security Level dialog box to configure the security settings.
  - Default Level—Sets the security level back to the default.
- IMSI Prefix Filtering—Opens the IMSI Prefix Filtering dialog box to configure IMSI prefix filters.
- Advanced View—Lets you configure the security settings.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

GTP Inspect Map Advanced View

Configuration > Global Objects > Inspect Maps > GTP > GTP Inspect Map > Advanced View

The GTP Inspect Map Advanced View pane lets you configure the settings for the inspect map.

Fields

- Name—Shows the name of the previously configured GTP map.
- Description—Enter the description of the GTP map, up to 200 characters in length.
- 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 security appliance 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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, protocol state tracking, H.323 call duration enforcement, and audio/video control.

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.
- 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

- Customize—Opens the Customize Security Level dialog box for additional settings.
- Default Level—Sets the security level back to the default level of Medium.
- Phone Number Filtering—Opens the Phone Number Filtering dialog box to configure phone number filters.

- H.323 Inspect Maps—Table that lists the defined H.323 inspect maps. The defined inspect maps are also listed in the H.323 area of the Inspect Maps tree.
- Add—Adds the new H.323 inspect map to the defined list in the H.323 Inspect Maps table and to the H.323 area of the Inspect Maps tree. To configure the new H.323 map, select the H.323 entry in Inspect Maps tree.
- Delete—Deletes the application inspection map selected in the H.323 Inspect Maps table and from the H.323 area of the Inspect Maps tree.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Customize Security Level

Configuration > Global Objects > Inspect Maps > H323 > Customize Security Level

The Customize Security Level dialog box lets you configure the security settings for previously configured H.323 application inspection maps.

Fields

- Settings—Specifies H.323 security settings and actions.
  - 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.
  - 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.
  - 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.
- Reset to predefined security level—Resets the security level settings to the predefined levels of high, medium, or low. Default is low.

- Reset to—Specifies high, medium, or low security setting.
- Reset—Reset settings to selected level.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

H.323 Inspect Map Basic View

Configuration > Global Objects > Inspect Maps > H323 > H323 Inspect Map > Basic View

The H323 Inspect Map Basic View pane shows the configured settings for the H323 inspect map. The Advanced View lets you configure the settings.

Fields

- Name—Shows the name of the previously configured H323 map.
- Description—Enter the description of the H323 map, up to 200 characters in length.
- Security Level—Shows the current security settings.
  - Customize—Opens the Customize Security Level dialog box to configure the security settings.
  - Default Level—Sets the security level back to the default.
- Phone Number Filtering—Opens the Phone Number Filtering dialog box which lets you configure the settings for a phone number filter.
- Advanced View—Lets you configure the security settings.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

H.323 Inspect Map Advanced View

Configuration > Global Objects > Inspect Maps > H323 > H323 Inspect Map > Advanced View

The H.323 Inspect Map Advanced View pane lets you configure the settings for the inspect map.

Fields

- Name—Shows the name of the previously configured H.323 map.
- Description—Enter the description of the H.323 map, up to 200 characters in length.
- 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.
- 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

HTTP Inspect Map

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

- 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.
- 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

– Customize—Opens the Customize Security Level dialog box for additional settings.
– Default Level—Sets the security level back to the default level of Medium.
– URI Filtering—Opens the URI Filtering dialog box to configure URI filters.

- HTTP Inspect Maps—Table that lists the defined HTTP inspect maps. The defined inspect maps are also listed in the HTTP area of the Inspect Maps tree.
- Add—Adds the new HTTP inspect map to the defined list in the HTTP Inspect Maps table and to the HTTP area of the Inspect Maps tree. To configure the new HTTP map, select the HTTP entry in Inspect Maps tree.
- Delete—Deletes the application inspection map selected in the HTTP Inspect Maps table and from the HTTP area of the Inspect Maps tree.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Customize Security Level

Configuration > Global Objects > Inspect Maps > HTTP > Customize Security Level

The Customize Security Level dialog box lets you configure the security settings for previously configured HTTP application inspection maps.

Fields

- Settings—Specifies HTTP security settings and actions.
  – Check for protocol violations—Checks for HTTP protocol violations.

Action—Drop Connection, Reset, Log.

Log—Enable or disable.

- Drop connections for unsafe methods—Checks for unsafe methods and drops the connection.

  Allow Only—GET and HEAD, GET, HEAD, and POST.

- Drop connections for requests with non-ASCII headers—Checks for non-ASCII characters in the message header.

- Reset to predefined security level—Resets the security level settings to the predefined levels of high, medium, or low. Default is low.
  - Reset to—Specifies high, medium, or low security setting.
  - Reset—Reset settings to selected level.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

URI Filtering

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

HTTP Inspect Map Basic View

Configuration > Global Objects > Inspect Maps > HTTP > HTTP Inspect Map > Basic View

The HTTP Inspect Map Basic View pane shows the configured settings for the HTTP inspect map. The Advanced View lets you configure the settings.

Fields

- Name—Shows the name of the previously configured HTTP map.
- Description—Enter the description of the HTTP map, up to 200 characters in length.
- Security Level—Shows the current security settings.
  - Customize—Opens the Customize Security Level dialog box to configure the security settings.
  - Default Level—Sets the security level back to the default.
- URI Filtering—Opens the URI Filtering dialog box which lets you configure the settings for an URI filter.
- Advanced View—Lets you configure the security settings.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

HTTP Inspect Map Advanced View

Configuration > Global Objects > Inspect Maps > HTTP > HTTP Inspect Map > Advanced View

The HTTP Inspect Map Advanced View pane lets you configure the settings for the inspect map.

Fields

- Name—Shows the name of the previously configured HTTP map.
- Description—Enter the description of the HTTP map, up to 200 characters in length.
- 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit HTTP Map

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Instant Messaging (IM) Inspect Map

Configuration > Global Objects > Inspect Maps > IM

The IM pane lets you view previously configured Instant Messaging (IM) application inspection maps. An Instant Messaging (IM) map lets you change the default configuration values used for Instant Messaging (IM) application inspection.

Instant Messaging (IM) application inspection provides detailed access control to control network usage. It also helps stop leakage of confidential data and propagations of network threats. A regular expression database search representing various patterns for Instant Messaging (IM) protocols to be filtered is applied. A syslog is generated if the flow is not recognized.

The scope can be limited by using an access list to specify any traffic streams to be inspected. For UDP messages, a corresponding UDP port number is also configurable. Inspection of Yahoo! Messenger and MSN Messenger instant messages are supported.

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.
- IM Inspect Maps—Table that lists the defined IM inspect maps. The defined inspect maps are also listed in the IM area of the Inspect Maps tree.
- Add—Adds the new IM inspect map to the defined list in the IM Inspect Maps table and to the IM area of the Inspect Maps tree. To configure the new IM map, select the IM entry in Inspect Maps tree.
- Delete—Deletes the application inspection map selected in the IM Inspect Maps table and from the IM area of the Inspect Maps tree.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Instant Messaging (IM) Inspect Map View

Configuration > Global Objects > Inspect Maps > IM > IM Inspect Map > View

The IM Inspect Map View pane lets you configure the settings for the inspect map.

Fields

- Name—Shows the name of the previously configured IM map.
- Description—Enter the description of the IM map, up to 200 characters in length.
- Match Type—Shows the match type, which can be a positive or negative match.
- Criterion—Shows the criterion of the IM inspection.
- Value—Shows the value to match in the IM inspection.
- Action—Shows the action if the match condition is met.
- Log—Shows the log state.
- Add—Opens the Add IM Inspect dialog box to add an IM inspection.
- Edit—Opens the Edit IM Inspect dialog box to edit an IM inspection.
- Delete—Deletes an IM inspection.

- Move Up—Moves an inspection up in the list.
- Move Down—Moves an inspection down in the list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit IM Map

Configuration > Global Objects > Inspect Maps > IM > IM Inspect Map > View > Add/Edit IM Inspect

The Add/Edit IM Inspect dialog box lets you define the match criterion and value for the IM inspect map.

Fields

- Single Match—Specifies that the IM 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 IM traffic to match.
  - Protocol—Match IM protocols.
  - Service—Match IM services.
  - Source IP Address—Match source IP address.
  - Destination IP Address—Match destination IP address.
  - Version—Match IM file transfer service version.
  - Client Login Name—Match client login name from IM service.
  - Client Peer Login Name—Match client peer login name from IM service.
  - Filename—Match filename form IM file transfer service.
- Protocol Criterion Values—Specifies which IM protocols to match.
  - Yahoo! Messenger—Specifies to match Yahoo! Messenger instant messages.
  - MSN Messenger—Specifies to match MSN Messenger instant messages.
- Service Criterion Values—Specifies which IM services to match.
  - Chat—Specifies to match IM message chat service.
  - Conference—Specifies to match IM conference service.
  - File Transfer—Specifies to match IM file transfer service.
  - Games—Specifies to match IM gaming service.
  - Voice Chat—Specifies to match IM voice chat service (not available for Yahoo IM)

  - Web Cam—Specifies to match IM webcam service.
- Source IP Address Criterion Values—Specifies to match the source IP address of the IM service.
  - IP Address—Enter the source IP address of the IM service.
  - IP Mask—Mask of the source IP address.
- Destination IP Address Criterion Values—Specifies to match the destination IP address of the IM service.
  - IP Address—Enter the destination IP address of the IM service.
  - IP Mask—Mask of the destination IP address.
- Version Criterion Values—Specifies to match the version from the IM file transfer service. 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.
- Client Login Name Criterion Values—Specifies to match the client login name from the IM service. 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.
- Client Peer Login Name Criterion Values—Specifies to match the client peer login name from the IM service. 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.
- Filename Criterion Values—Specifies to match the filename from the IM file transfer service. 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.
- Multiple Matches—Specifies multiple matches for the IM inspection.

  - IM Traffic Class—Specifies the IM traffic class match.
  - Manage—Opens the Manage IM Class Maps dialog box to add, edit, or delete IM Class Maps.
- Action—Drop connection, reset, or log.
- Log—Enable or disable.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

IPSec Pass Through Inspect Map

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 access list.

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.
- 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 Customize Security Level dialog box for additional settings.
  - Default Level—Sets the security level back to the default level of Low.
- IPSec Pass Through Inspect Maps—Table that lists the defined IPSec Pass Through inspect maps. The defined inspect maps are also listed in the IPSec Pass Through area of the Inspect Maps tree.

- Add—Adds the new IPSec Pass Through inspect map to the defined list in the IPSec Pass Through Inspect Maps table and to the IPSec Pass Through area of the Inspect Maps tree. To configure the new IPSec Pass Through map, select the IPSec Pass Through entry in Inspect Maps tree.
- Delete—Deletes the application inspection map selected in the IPSec Pass Through Inspect Maps table and from the IPSec Pass Through area of the Inspect Maps tree.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Customize Security Level

Configuration > Global Objects > Inspect Maps > IPSec Pass Through> Customize Security Level

The Customize Security Level dialog box lets you configure the security settings for previously configured IPSec Pass Through application inspection maps.

Fields

- Settings—Specifies IPSec Pass Through security settings and actions.
  - 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.
- Reset to predefined security level—Resets the security level settings to the predefined levels of high, medium, or low. Default is low.
  - Reset to—Specifies high, medium, or low security setting.
  - Reset—Reset settings to selected level.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

IPSec Pass Through Inspect Map Basic View

Configuration > Global Objects > Inspect Maps > IPSec Pass Through>
IPSec Pass Through Inspect Map > Basic View

The IPSec Pass Through Inspect Map Basic View pane lets you configure basic settings for the inspect map.

Fields

- Description—Enter the description of the inspect map, up to 200 characters in length.
- 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 Customize Security Level dialog box for additional settings.
  - Default Level—Sets the security level back to the default level of Low.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

IPSec Pass Through Inspect Map Advanced View

Configuration > Global Objects > Inspect Maps >IPSec Pass Through > IPSec Pass Through Inspect Map > Advanced View

The IPSec Pass Through Inspect Map Advanced View pane lets you configure advanced settings for the inspect map.

Fields

- Name—Shows the name of the previously configured IPSec Pass Through map.
- Description—Enter the description of the IPSec Pass Through map, up to 200 characters in length.
- 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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

- 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.
- Command Queue Size—Specifies the maximum number of commands to queue. The valid range is from 1 to 2147483647.
- Gateways and Call Agents—Opens the Gateways and Call Agents dialog box to add an MGCP map.
- MGCP Inspect Maps—Table that lists the defined MGCP inspect maps. The defined inspect maps are also listed in the MGCP area of the Inspect Maps tree.

- Add—Adds the new MGCP inspect map to the defined list in the MGCP Inspect Maps table and to the MGCP area of the Inspect Maps tree. To configure the new MGCP map, select the MGCP entry in Inspect Maps tree.
- Delete—Deletes the application inspection map selected in the MGCP Inspect Maps table and from the MGCP area of the Inspect Maps tree.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

MGCP Inspect Map View

Configuration > Global Objects > Inspect Maps > MGCP > MGCP Inspect Map > View

The MGCP Inspect Map View pane lets you configure the settings for the inspect map.

Fields

- Name—Shows the name of the previously configured MGCP map.
- 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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.
- Gateways area
  - 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

NetBIOS Inspect Map

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

- 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.
- Check for protocol violations—Checks for protocol violations and executes specified action.
  - Action—Drop packet or log.
  - Log—Enable or disable.
- NetBIOS Inspect Maps—Table that lists the defined NetBIOS inspect maps. The defined inspect maps are also listed in the NetBIOS area of the Inspect Maps tree.
- Add—Adds the new NetBIOS inspect map to the defined list in the NetBIOS Inspect Maps table and to the NetBIOS area of the Inspect Maps tree. To configure the new NetBIOS map, select the NetBIOS entry in Inspect Maps tree.
- Delete—Deletes the application inspection map selected in the NetBIOS Inspect Maps table and from the NetBIOS area of the Inspect Maps tree.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

NetBIOS Inspect Map View

Configuration > Global Objects > Inspect Maps > NetBIOS > NetBIOS Inspect Map > View

The NetBIOS Inspect Map View pane lets you configure the settings for the inspect map.

Fields

- Name—Shows the name of the previously configured NetBIOS map.
- 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

RADIUS Inspect Map

Configuration > Global Objects > Inspect Maps > RADIUS

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. ou 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

RADIUS Inspect Map Host

Configuration > Global Objects > Inspect Maps > RADIUS > 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

RADIUS Inspect Map Other

Configuration > Global Objects > Inspect Maps > RADIUS > 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.
  - 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).
- Enable detection of GPRS accounting—Enables detection of GPRS accounting. This option is only available when GTP/GPRS license is enabled.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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

- 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.
- 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.

  - Customize—Opens the Customize Security Level dialog box for additional settings.
  - Default Level—Sets the security level back to the default level of Low.

- Message ID Filtering—Opens the Messaging ID Filtering dialog box for configuring message ID filters.
- SCCP (Skinny) Inspect Maps—Table that lists the defined SCCP (Skinny) inspect maps. The defined inspect maps are also listed in the SCCP (Skinny) area of the Inspect Maps tree.
- Add—Adds the new SCCP (Skinny) inspect map to the defined list in the SCCP (Skinny) Inspect Maps table and to the SCCP (Skinny) area of the Inspect Maps tree. To configure the new SCCP (Skinny) map, select the SCCP (Skinny) entry in Inspect Maps tree.
- Delete—Deletes the application inspection map selected in the SCCP (Skinny) Inspect Maps table and from the SCCP (Skinny) area of the Inspect Maps tree.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Customize Security Level

Configuration > Global Objects > Inspect Maps > SCCP (Skinny) > Customize Security Level

The Customize Security Level dialog box lets you configure the security settings for previously configured SCCP (Skinny) application inspection maps.

Fields

- Settings—Specifies SCCP (Skinny) security settings and actions.
  - 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 (0x0 to 0xffff).
  - SCCP Prefix Length—Specifies prefix length value in Skinny messages (4 to 4,294,967,295).

    Minimum Prefix Length—Specify minimum value of SCCP prefix length allowed.

    Maximum Prefix Length—Specify maximum value of SCCP prefix length allowed.
  - Enable media timeout—Enables media timeout.

    Media Timeout—Specify timeout value for media connections (0:0:01 to 1993:0:0).
  - Enable signaling timeout—Enables signaling timeout.

Signaling Timeout—Specify timeout value for signaling connections (0:0:01 to 1993:0:0).

– 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.

- Reset to predefined security level—Resets the security level settings to the predefined levels of high, medium, or low. Default is low.
  – Reset to—Specifies high, medium, or low security setting.
  – Reset—Reset settings to selected level.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

SCCP (Skinny) Inspect Map Basic View

Configuration > Global Objects > Inspect Maps > SCCP (Skinny) > SCCP (Skinny) Inspect Map > Basic View

The SCCP (Skinny) Inspect Map Basic View pane shows the configured settings for the SCCP (Skinny) inspect map. The Advanced View lets you configure the settings.

Fields

- Name—Shows the name of the previously configured SCCP (Skinny) map.
- Description—Enter the description of the DNS map, up to 200 characters in length.
- Security Level—Shows the current security settings.
  - Customize—Opens the Customize Security Level dialog box to configure the security settings.
  - Default Level—Sets the security level back to the default.
  - Message ID Filtering—Opens the Messaging ID Filtering dialog box for configuring message ID filters.
- Advanced View—Lets you configure the security settings.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

SCCP (Skinny) Inspect Map Advanced View

Configuration > Global Objects > Inspect Maps > SCCP (Skinny) > SCCP (Skinny) Inspect Map > Advanced View

The SCCP (Skinny) Inspect Map Advanced View pane lets you configure the inspect map settings.

Fields

- Name—Shows the name of the previously configured SCCP (Skinny) map.
- Description—Enter the description of the DNS map, up to 200 characters in length.
- 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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

- 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.
- 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 Customize Security Level dialog box for additional settings.
– Default Level—Sets the security level back to the default level of Low.

- SIP Inspect Maps—Table that lists the defined SIP inspect maps. The defined inspect maps are also listed in the SIP area of the Inspect Maps tree.
- Add—Adds the new SIP inspect map to the defined list in the SIP Inspect Maps table and to the SIP area of the Inspect Maps tree. To configure the new SIP map, select the SIP entry in Inspect Maps tree.
- Delete—Deletes the application inspection map selected in the SIP Inspect Maps table and from the SIP area of the Inspect Maps tree.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Customize Security Level

Configuration > Global Objects > Inspect Maps > SIP > Customize Security Level

The Customize Security Level dialog box lets you configure the security settings for previously configured SIP application inspection maps.

Fields

- Settings—Lets you configure additional SIP settings, including RTP and SIP conformance.
  - 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.
  - Hide server's and endpoint's IP addresses—Enables IP address privacy. Disabled by default.
  - Mask software version and non-SIP URIs—Enables non-SIP URI inspection in Alert-Info and Call-Info headers.
  - Ensure that number of hops to destination is greater than 0—Enables check for the value of Max-Forwards header is zero.
- RTP Conformance
  - 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
  - Do not perform state checking and header validation—Disables SIP state checking.
  - Drop packets that fail state checking—Drops packets that fail state checking.
  - Drop connections that fail state checking and packets that fail header validation—Drops connections that fail state checking and packets that fail header validation of SIP messages.
- Reset to Predefined Security Level—Resets the security level settings to the predefined levels of high, medium, or low.
  - Reset To—Resets the security level to high, medium, or low.
- Reset—Resets all security settings to the default. The default pinhole timeout is one minute. The default endpoint mapper settings are none.Criterion—Specifies which criterion of SIP traffic to match.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

SIP Inspect Map Basic View

Configuration > Global Objects > Inspect Maps > SIP > SIP Inspect Map > Basic View

The SIP Inspect Map Basic View pane shows the configured settings for the SIP inspect map. The Advanced View lets you configure the settings.

Fields

- Name—Shows the name of the previously configured SIP map.
- Description—Enter the description of the DNS map, up to 200 characters in length.
- Security Level—Shows the current security settings.
  - Customize—Opens the Customize Security Level dialog box to configure the security settings.
  - Default Level—Sets the security level back to the default.
- Advanced View—Lets you configure the security settings.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

SIP Inspect Map Advanced View

Configuration > Global Objects > Inspect Maps > SIP > SIP Inspect Map > Advanced View

The SIP Inspect Map Advanced View pane lets you configure the inspect map settings.

Fields

- Name—Shows the name of the previously configured SIP map.
- Description—Enter the description of the DNS map, up to 200 characters in length.
- 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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. Check a map and click **Edit** to view or change an existing map.
- Disallowed SNMP Versions—Identifies the SNMP versions that have been disallowed for a specific SNMP application inspection map.
- Add—Displays the Add SNMP dialog box, which you can use to define a new application inspection map.
- Edit—Displays the Edit SNMP 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Configuring Regular Expressions

This section describes how to configure regular expressions, and includes the following topics:

Regular Expressions

Configuration > Global Objects > Regular Expressions

Some Configuring Class Maps and Configuring Inspect Maps can specify regular expressions to match text inside a packet. Be sure to create the regular expressions before you configure the class map or inspect map, either singly or grouped together in a regular expression class map.

A regular expression matches text strings either literally as an exact string, or by using *metacharacters* so you can match multiple variants of a text string. You can use a regular expression to match the content of certain application traffic; for example, you can match body text inside an HTTP packet.

Fields

- Regular Expressions—Shows the regular expressions
  - Name—Shows the regular expression names.
  - Value—Shows the regular expression definitions.
  - Add—Adds a regular expression.
  - Edit—Edits a regular expression.
  - Delete—Deletes a regular expression.
- Regular Expression Classes—Shows the regular expression class maps.
  - Name—Shows the regular expression class map name.
  - Match Conditions—Shows the match type and regular expressions in the class map.

    Match Type—Shows the match type, which for regular expressions is always a positive match type (shown by the icon with the equal sign (=)) the criteria. (Inspection class maps allow you to create negative matches as well (shown by the icon with the red circle)). If more than one regular expression is in the class map, then each match type icon appears with "OR" next it, to indicate that this class map is a "match any" class map; traffic matches the class map if only one regular expression is matched.

    Regular Expression—Lists the regular expressions included in each class map.
  - Description—Shows the description of the class map.
  - Add—Adds a regular expression class map.
  - Edit—Edits a regular expression class map.
  - Delete—Deletes a regular expression class map.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Regular Expression

Configuration > Global Objects > Regular Expressions > Add/Edit a Regular Expression

The Add/Edit Regular Expression dialog box lets you define and test a regular expression.

Fields

- Name—Enter the name of the regular expression, up to 40 characters in length.
- Value—Enter the regular expression, up to 100 characters in length. You can enter the text manually, using the metacharacters in Table 6-1, or you can click **Build** to use the Build Regular Expression dialog box.

  Table 6-1 lists the metacharacters that have special meanings.

Table 6-1 regex Metacharacters

Character	Description	Notes
.	Dot	Matches any single character. For example, **d.g** matches dog, dag, dtg, and any word that contains those characters, such as doggonnit.
(*exp*)	Subexpression	A subexpression segregates characters from surrounding characters, so that you can use other metacharacters on the subexpression. For example, **d(o|a)g** matches dog and dag, but **do|ag** matches do and ag. A subexpression can also be used with repeat quantifiers to differentiate the characters meant for repetition. For example, **ab(xy){3}z** matches abxyxyxyz.
|	Alternation	Matches either expression it separates. For example, **dog|cat** matches dog or cat.
?	Question mark	A quantifier that indicates that there are 0 or 1 of the previous expression. For example, **lo?se** matches lse or lose. **Note** You must enter **Ctrl+V** and then the question mark or else the help function is invoked.
*	Asterisk	A quantifier that indicates that there are 0, 1 or any number of the previous expression. For example, **lo*se** matches lse, lose, loose, etc.
+	Plus	A quantifier that indicates that there is at least 1 of the previous expression. For example, **lo+se** matches lose and loose, but not lse.
{*x*}	Repeat quantifier	Repeat exactly *x* times. For example, **ab(xy){3}z** matches abxyxyxyz.
{*x*,}	Minimum repeat quantifier	Repeat at least *x* times. For example, **ab(xy){2,}z** matches abxyxyz, abxyxyxyz, etc.
[*abc*]	Character class	Matches any character in the brackets. For example, **[abc]** matches a, b, or c.
[^*abc*]	Negated character class	Matches a single character that is not contained within the brackets. For example, **[^abc]** matches any character other than a, b, or c. **[^A-Z]** matches any single character that is not an uppercase letter.
[*a*-*c*]	Character range class	Matches any character in the range. **[a-z]** matches any lowercase letter. You can mix characters and ranges: **[abcq-z]** matches a, b, c, q, r, s, t, u, v, w, x, y, z, and so does [**a-cq-z**]. The dash (-) character is literal only if it is the last or the first character within the brackets: **[abc-]** or **[-abc]**.
""	Quotation marks	Preserves trailing or leading spaces in the string. For example, **" test"** preserves the leading space when it looks for a match.
^	Caret	Specifies the beginning of a line.

Table 6-1 regex Metacharacters (continued)

Character	Description	Notes
\	Escape character	When used with a metacharacter, matches a literal character. For example, \[matches the left square bracket.
char	Character	When character is not a metacharacter, matches the literal character.
\r	Carriage return	Matches a carriage return 0x0d.
\n	Newline	Matches a new line 0x0a.
\t	Tab	Matches a tab 0x09.
\f	Formfeed	Matches a form feed 0x0c.
\x*NN*	Escaped hexadecimal number	Matches an ASCII character using hexadecimal (exactly two digits).
NNN	Escaped octal number	Matches an ASCII character as octal (exactly three digits). For example, the character 040 represents a space.

- Build—Helps you build a regular expression using the Build Regular Expression dialog box.
- Test—Tests a regular expression against some sample text.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Build Regular Expression

Configuration > Global Objects > Regular Expressions > Add/Edit a Regular Expression > Build Regular Expression

The Build Regular Expression dialog box lets you construct a regular expression out of characters and metacharacters. Fields that insert metacharacters include the metacharacter in parentheses in the field name.

Fields

Build Snippet—This area lets you build text snippets of regular text or lets you insert a metacharacter into the Regular Expression field.

- Starts at the beginning of the line (^)—Indicates that the snippet should start at the beginning of a line, using the caret (^) metacharacter. Be sure to insert any snippet with this option at the beginning of the regular expression.
- Specify Character String—Enter a text string manually.

  - Character String—Enter a text string.
  - Escape Special Characters—If you entered any metacharacters in your text string that you want to be used literally, check this box to add the backslash (\) escape character before them. for example, if you enter “example.com,” this option converts it to “example\.com”.
  - Ignore Case—If you want to match upper and lower case characters, this check box automatically adds text to match both upper and lower case. For example, entering “cats” is converted to “[cC][aA][tT][sS]”.

- Specify Character—Lets you specify a metacharacter to insert in the regular expression.
  - Negate the character—Specifies not to match the character you identify.
  - Any character (.)—Inserts the period (.) metacharacter to match any character. For example, **d.g** matches dog, dag, dtg, and any word that contains those characters, such as doggonnit.
  - Character set—Inserts a character set. Text can match any character in the set. Sets include:

    [0-9A-Za-z]

    [0-9]

    [A-Z]

    [a-z]

    [aeiou]

    [\n\f\r\t] (which matches a new line, form feed, carriage return, or a tab)

    For example, if you specify [0-9A-Za-z], then this snippet will match any character from A to Z (upper or lower case) or any digit 0 through 9.
  - Special character—Inserts a character that requires an escape, including \, ?, *, +, |, ., [, (, or ^. The escape character is the backslash (\), which is auatomatically entered when you choose this option.
  - Whitespace character—Whitespace characters include \n (new line), \f (form feed), \r (carriage return), or \t (tab).
  - Three digit octal number—Matches an ASCII character as octal (up to three digits). For example, the character \040 represents a space. The backslash (\) is entered automatically.
  - Two digit hexadecimal number—Matches an ASCII character using hexadecimal (exactly two digits). The backslash (\) is entered automatically.
  - Specified character—Enter any single character.
- Snippet Preview—*Display only.* Shows the snippet as it will be entered in the regular expression.
- Append Snippet—Adds the snippet to the end of the regular expression.
- Append Snippet as Alternate—Adds the snippet to the end of the regular expression separated by a pipe (|), which matches either expression it separates. For example, **dog|cat** matches dog or cat.
- Insert Snippet at Cursor—Inserts the snippet at the cursor.

Regular Expression—This area includes regular expression text that you can enter manually and build with snippets. You can then select text in the Regular Expression field and apply a quantifier to the selection.

- Selection Occurrances—Select text in the Regular Expression field, click one of the following options, and then click **Apply to Selection**. For example, if the regular expression is “test me,” and you select “me” and apply **One or more times**, then the regular expression changes to “test (me)+”.
  - Zero or one times (?)—A quantifier that indicates that there are 0 or 1 of the previous expression. For example, **lo?se** matches lse or lose.

  - One or more times (+)—A quantifier that indicates that there is at least 1 of the previous expression. For example, **lo+se** matches lose and loose, but not lse.
  - Any number of times (*)—A quantifier that indicates that there are 0, 1 or any number of the previous expression. For example, **lo*se** matches lse, lose, loose, etc.
  - At least—Repeat at least *x* times. For example, **ab(xy){2,}z** matches abxyxyz, abxyxyxyz, etc.
  - Exactly—Repeat exactly *x* times. For example, **ab(xy){3}z** matches abxyxyxyz.
  - Apply to Selection—Applies the quantifier to the selection.
- Test—Tests a regular expression against some sample text.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Test Regular Expression

Configuration > Global Objects > Regular Expressions > Add/Edit a Regular Expression > Test Regular Expression

The Test Regular Expression dialog box lets you test input text against a regular expression to make sure it matches as you intended.

Fields

- Regular Expression—Enter ther regular expression you want to test. By default, the regular expression you entered in the Add/Edit Regular Expression or Build Regular Expression dialog box is input into this field. If you change the regular expression during your testing, and click **OK**, the changes are inherited by the Add/Edit Regular Expression or Build Regular Expression dialog boxes. Click **Cancel** to dismiss your changes.
- Test String—Enter a text string that you expect to match the regular expression.
- Test—Tests the Text String against the Regular Expression,
- Test Result—*Display only.* Shows if the test succeeded or failed.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Regular Expression Class Map

Configuration > Global Objects > Regular Expressions > Add/Edit Regular Expression Class Map

The Add/Edit Regular Expression Class Map dialog box groups regular expressions together. A regular expression class map can be used by inspection class maps and inspection policy maps.

Fields

- Name—Enter a name for the class map, up to 40 characters in length. The name "class-default" is reserved. All types of class maps use the same name space, so you cannot reuse a name already used by another type of class map.
- Description—Enter a description, up to 200 characters in length.
- Available Regular Expressions—Lists the regular expressions that are not yet assigned to the class map.
  - Edit—Edits the selected regular expression.
  - New—Creates a new regular expression.
- Add—Adds the selected regular expression to the class map.
- Remove—Removes the selected regular expression from the class map.
- Configured Match Conditions—Shows the regular expressions in this class map, along with the match type.
  - Match Type—Shows the match type, which for regular expressions is always a positive match type (shown by the icon with the equal sign (=)) the criteria. (Inspection class maps allow you to create negative matches as well (shown by the icon with the red circle)). If more than one regular expression is in the class map, then each match type icon appears with "OR" next it, to indicate that this class map is a "match any" class map; traffic matches the class map if only one regular expression is matched.
  - Regular Expression—Lists the regular expression names in this class map.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

TCP Maps

Configuration > Global Objects > TCP Maps (You can get to this pane through various paths.)

Use the TCP Maps option to create a reusable component that defines the TCP connection settings for different traffic flows. After creating a TCP map, you can associate these connection settings with traffic of a specific type using a security policy. You use the Service Policy Rules option on the Security Policy pane to define the traffic criteria and to associate the service policy rule with a specific interface or to apply it to all the interfaces on the security appliance.

The TCP Maps pane lets you customize inspection on TCP flow for both through and to the box traffic.

Fields

- Map Name—Lists a TCP map name used to apply a TCP map.
- Urgent Flag—Lists whether the URG pointer is cleared or allowed through the security appliance.
- Window Variation—Lists whether a connection that has changed its window size unexpectedly is allowed or dropped.
- SYN Data—Lists whether SYN packets with data are allowed or dropped.
- TTL Evasion Protection—Lists whether the TTL evasion protection offered by the security appliance is enabled or disabled.
- Exceed MSS—Lists whether packets that exceed MSS set by peer are allowed or dropped.
- Check Retransmission—Lists whether the retransmit data check is enabled or disabled.
- Verify Checksum—Lists whether checksum verification is enabled or disabled.
- Reserved Bits—Lists the status of the reserved flags policy.
- TCP Option—Lists the behavior of packets with TCP option value configured. The default action is to clear the options and allow the packets.
  - Clear Selective Ack—Lists whether the selective-ack TCP option is allowed or cleared.
  - Clear TCP Timestamp—Lists whether the TCP timestamp option is allowed or cleared.
  - Clear Window Scale—Lists whether the window scale timestamp option is allowed or cleared.
  - Range—Lists 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.
- Queue Size—Lists the maximum number of out-of-order packets that can be queued for a TCP connection. Default is 0.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit TCP Map

Configuration > Global Objects > TCP Maps > Add/Edit TCP Map (You can get to this dialog box through various paths.)

The Add/Edit TCP Maps dialog box lets you define the class of traffic and customize the TCP inspection with TCP maps. Apply the TCP map using policy map and activate TCP inspection using service policy.

Fields

- TCP Map Name—Specifies a TCP map name to use to apply a TCP map.
- Clear Urgent Flag—Allows or clears the URG pointer through the security appliance.

- Drop SYN Packets With Data—Allows or drops SYN packets with data.
- Drop Connection on Window Variation—Drops a connection that has changed its window size unexpectedly.
- Enable TTL Evasion Protection—Enables or disables the TTL evasion protection offered by the security appliance.
- Drop Packets that Exceed Maximum Segment Size—Allows or drops packets that exceed MSS set by peer.
- Verify TCP Checksum—Enables and disables checksum verification.
- Check if transmitted data is the same as original—Enables and disables the retransmit data checks.
- Reserved Bits—Sets the reserved flags policy in the security appliance.
  - Clear and allow
  - Allow only
  - Drop
- TCP Option—Configure the behavior of packets with TCP option value configured. The default action is to clear the options and allow the packets.
  - Clear Selective Ack—Allows or clears the selective-ack TCP options.
  - Clear TCP Timestamp—Allows or clears the TCP timestamp option.
  - Clear Window Scale—Allows or clears the window scale timestamp option.
- Range—Valid TCP options ranges should fall within 6-7 and 9-255. The lower bound should be less than or equal to the upper bound.
- Action—Allow or drop.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Configuring Time Ranges

Configuration > Global Objects > Time Ranges (You can get to this pane through multiple paths.)

Use the Time Ranges option to create a reusable component that defines starting and ending times that can be applied to various security features. Once you have defined a time range, you can select the time range and apply it to different options that require scheduling.

The time range feature lets you define a time range that you can attach to traffic rules, or an action. For example, you can attach an access list to a time range to restrict access to the security appliance.

A time range consists of a start time, an end time, and optional periodic entries.

Note Creating a time range does not restrict access to the device. This pane defines the time range only.

Fields

- Name—Specifies the name of the time range.
- Start Time—Specifies when the time range begins.
- End Time—Specifies when the time range ends.
- Periodic Entries—Specifies further constraints of active time of the range within the start and stop time specified.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Time Range

Configuration > Global Objects > Time Ranges > Add/Edit Time Range (You can get to this dialog box through multiple paths.)

The Add/Edit Time Range pane lets you define specific times and dates that you can attach to an action. For example, you can attach an access list to a time range to restrict access to the security appliance. The time range relies on the system clock of the security appliance; however, the feature works best with NTP synchronization.

Note Creating a time range does not restrict access to the device. This pane defines the time range only.

Fields

- Time Range Name—Specifies the name of the time range. The name cannot contain a space or quotation mark, and must begin with a letter or number.
- Start now/Started—Specifies either that the time range begin immediately or that the time range has begun already. The button label changes based on the Add/Edit state of the time range configuration. If you are adding a new time range, the button displays “Start Now.” If you are editing a time range for which a fixed start time has already been defined, the button displays “Start Now.” When editing a time range for which there is no fixed start time, the button displays “Started.”
- Start at—Specifies when the time range begins.
  - Month—Specifies the month, in the range of January through December.
  - Day—Specifies the day, in the range of 01 through 31.
  - Year—Specifies the year, in the range of 1993 through 2035.
  - Hour—Specifies the hour, in the range of 00 through 23.

  - Minute—Specifies the minute, in the range of 00 through 59.
- Never end—Specifies that there is no end to the time range.
- End at (inclusive)—Specifies when the time range ends. The end time specified is inclusive. For example, if you specified that the time range expire at 11:30, the time range is active through 11:30 and 59 seconds. In this case, the time range expires when 11:31 begins.
  - Month—Specifies the month, in the range of January through December.
  - Day—Specifies the day, in the range of 01 through 31.
  - Year—Specifies the year, in the range of 1993 through 2035.
  - Hour—Specifies the hour, in the range of 00 through 23.
  - Minute—Specifies the minute, in the range of 00 through 59.
- Periodic Time Ranges—Configures daily or weekly time ranges.
  - Add—Adds a periodic time range.
  - Edit—Edits the selected periodic time range.
  - Delete—Deletes the selected periodic time range.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Periodic Time Range

Configuration > Global Objects > Time Ranges > Add/Edit Time Range >
Add/Edit Periodic Time Range (You can get to this dialog box through multiple paths.)

The Add/Edit Periodic Time Range pane lets you fine time ranges further by letting you configure them on a daily or weekly basis.

Note Creating a time range does not restrict access to the device. This pane defines the time range only.

Fields

- Days of the week
  - Every day—Specifies every day of the week.
  - Weekdays—Specifies Monday through Friday.
  - Weekends—Specifies Saturday and Sunday.
  - On these days of the week—Lets you choose specific days of the week.
  - Daily Start Time—Specifies the hour and the minute that the time range begins.

  - Daily End Time (inclusive) area—Specifies the hour and the minute that the time range ends. The end time specified is inclusive.
- Weekly Interval
  - From—Lists the day of the week, Monday through Sunday.
  - Through—Lists the day of the week, Monday through Sunday.
  - Hour—Lists the hour, in the range of 00 through 23.
  - Minute—Lists the minute, in the range of 00 through 59.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Configuring Security Contexts

This chapter describes how to use security contexts and enable multiple context mode. This chapter includes the following sections:

Security Context Overview

You can partition a single security appliance into multiple virtual devices, known as security contexts. Each context is an independent device, with its own security policy, interfaces, and administrators. Multiple contexts are similar to having multiple standalone devices. Many features are supported in multiple context mode, including routing tables, firewall features, IPS, and management. Some features are not supported, including VPN and dynamic routing protocols.

In multiple context mode, the security appliance includes a configuration for each context that identifies the security policy, interfaces, and almost all the options you can configure on a standalone device. The system administrator adds and manages contexts by configuring them in the system configuration, which, like a single mode configuration, is the startup configuration. The system configuration identifies basic settings for the security appliance. The system configuration does not include any network interfaces or network settings for itself; rather, when the system needs to access network resources (such as downloading the contexts from the server), it uses one of the contexts that is designated as the admin context.

The admin context is just like any other context, except that when a user logs in to the admin context, then that user has system administrator rights and can access the system and all other contexts.

This section provides an overview of security contexts, and includes the following topics:

Common Uses for Security Contexts

You might want to use multiple security contexts in the following situations:

- You are a service provider and want to sell security services to many customers. By enabling multiple security contexts on the security appliance, you can implement a cost-effective, space-saving solution that keeps all customer traffic separate and secure, and also eases configuration.
- You are a large enterprise or a college campus and want to keep departments completely separate.
- You are an enterprise that wants to provide distinct security policies to different departments.
- You have any network that requires more than one security appliance.

Unsupported Features

Multiple context mode does not support the following features:

- Dynamic routing protocols

  Security contexts support only static routes. You cannot enable OSPF or RIP in multiple context mode.
- VPN
- Multicast

Context Configuration Files

Each context has its own configuration file that identifies the security policy, interfaces, and, for supported features, all the options you can configure on a standalone device. You can store context configurations on the internal Flash memory or the external Flash memory card, or you can download them from a TFTP, FTP, or HTTP(S) server.

In addition to individual security contexts, the security appliance also includes a system configuration that identifies basic settings for the security appliance, including a list of contexts. Like the single mode configuration, this configuration resides as the startup configuration.

The system configuration does not include any network interfaces or network settings for itself; rather, when the system needs to access network resources (such as downloading the contexts from a server), it uses one of the contexts that is designated as the admin context. The system configuration does include a specialized failover interface for failover traffic only. If your system is already in multiple context mode, or if you convert from single mode, the admin context is created automatically as a file on the internal Flash memory called admin.cfg. This context is named “admin.” If you do not want to use admin.cfg as the admin context, you can change the admin context.

How the Security Appliance Classifies Packets

Each packet that enters the security appliance must be classified, so that the security appliance can determine to which context to send a packet. This section includes the following topics:

- Valid Classifier Criteria, page 7-3
- Invalid Classifier Criteria, page 7-4

- Classification Examples, page 7-4

Note If the destination MAC address is a multicast or broadcast MAC address, the packet is duplicated and delivered to each context.

Valid Classifier Criteria

This section describes the criteria used by the classifier, and includes the following topics:

- Unique Interfaces, page 7-3
- Unique MAC Addresses, page 7-3
- NAT Configuration, page 7-3

Unique Interfaces

If only one context is associated with the ingress interface, the security appliance classifies the packet into that context. In transparent firewall mode, unique interfaces for contexts are required, so this method is used to classify packets at all times.

Unique MAC Addresses

If multiple contexts share an interface, then the classifier uses the interface MAC address. The security appliance lets you assign a different MAC address in each context to the same shared interface, whether it is a shared physical interface or a shared subinterface. By default, shared interfaces do not have unique MAC addresses; the interface uses the physical interface burned-in MAC address in every context. An upstream router cannot route directly to a context without unique MAC addresses. You can set the MAC addresses manually when you configure each interface (see Add/Edit Interface > Advanced), or you can automatically generate MAC addresses (see Security Contexts).

NAT Configuration

If you do not have unique MAC addresses, then the classifier intercepts the packet and performs a destination IP address lookup. All other fields are ignored; only the destination IP address is used. To use the destination address for classification, the classifier must have knowledge about the subnets located behind each security context. The classifier relies on the NAT configuration to determine the subnets in each context. The classifier matches the destination IP address to either a **static** command or a **global** command. In the case of the **global** command, the classifier does not need a matching **nat** command or an active NAT session to classify the packet. Whether the packet can communicate with the destination IP address after classification depends on how you configure NAT and NAT control.

For example, the classifier gains knowledge about subnets 10.10.10.0, 10.20.10.0 and 10.30.10.0 when the context administrators configure **static** commands in each context:

- Context A:

```
static (inside,shared) 10.10.10.0 10.10.10.0 netmask 255.255.255.0
```

- Context B:

```
static (inside,shared) 10.20.10.0 10.20.10.0 netmask 255.255.255.0
```

- Context C:

```
static (inside,shared) 10.30.10.0 10.30.10.0 netmask 255.255.255.0
```

Note For management traffic destined for an interface, the interface IP address is used for classification.

Invalid Classifier Criteria

The following configurations are not used for packet classification:

- NAT exemption—The classifier does not use a NAT exemption configuration for classification purposes because NAT exemption does not identify a mapped interface.
- Routing table—If a context includes a static route that points to an external router as the next-hop to a subnet, and a different context includes a **static** command for the same subnet, then the classifier uses the **static** command to classify packets destined for that subnet and ignores the static route.

Classification Examples

Figure 7-1 shows multiple contexts sharing an outside interface. The classifier assigns the packet to Context B because Context B includes the MAC address to which the router sends the packet.

Figure 7-1 Packet Classification with a Shared Interface using MAC Addresses

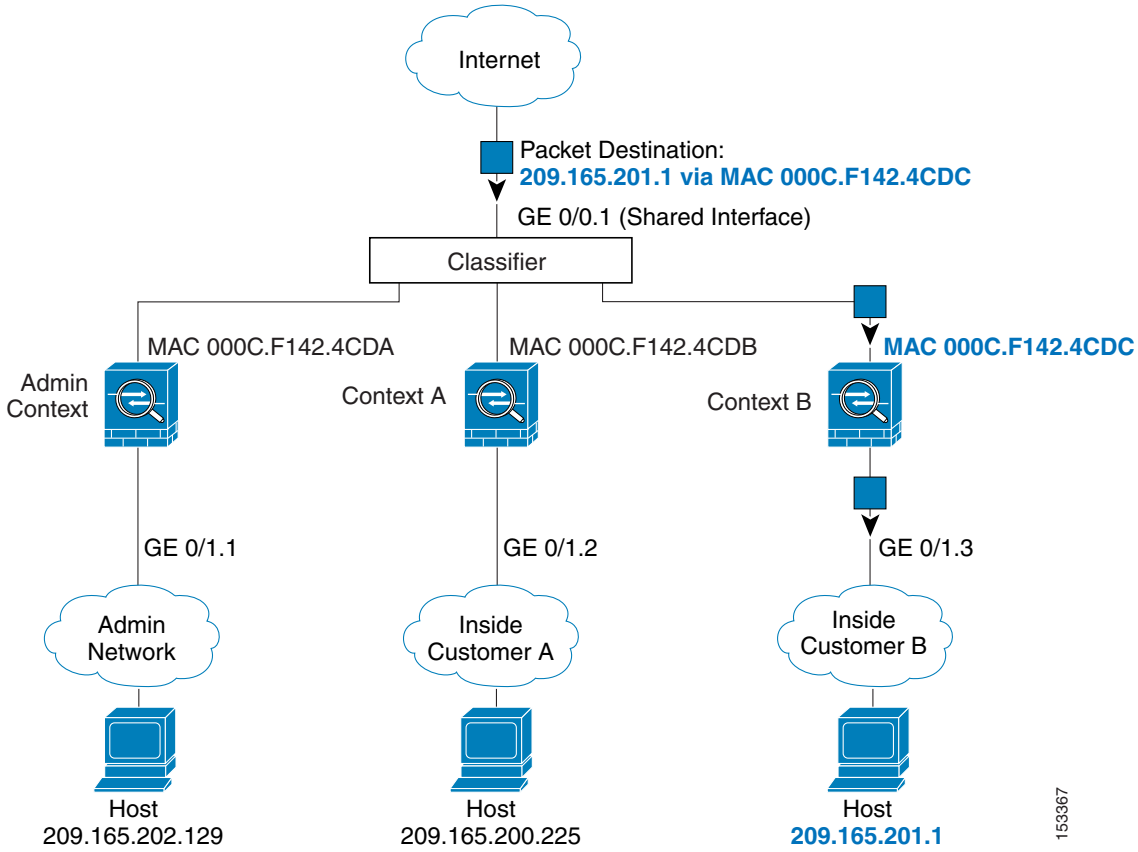

Figure 7-2 shows multiple contexts sharing an outside interface without MAC addresses assigned. The classifier assigns the packet to Context B because Context B includes the address translation that matches the destination address.

Figure 7-2 Packet Classification with a Shared Interface using NAT

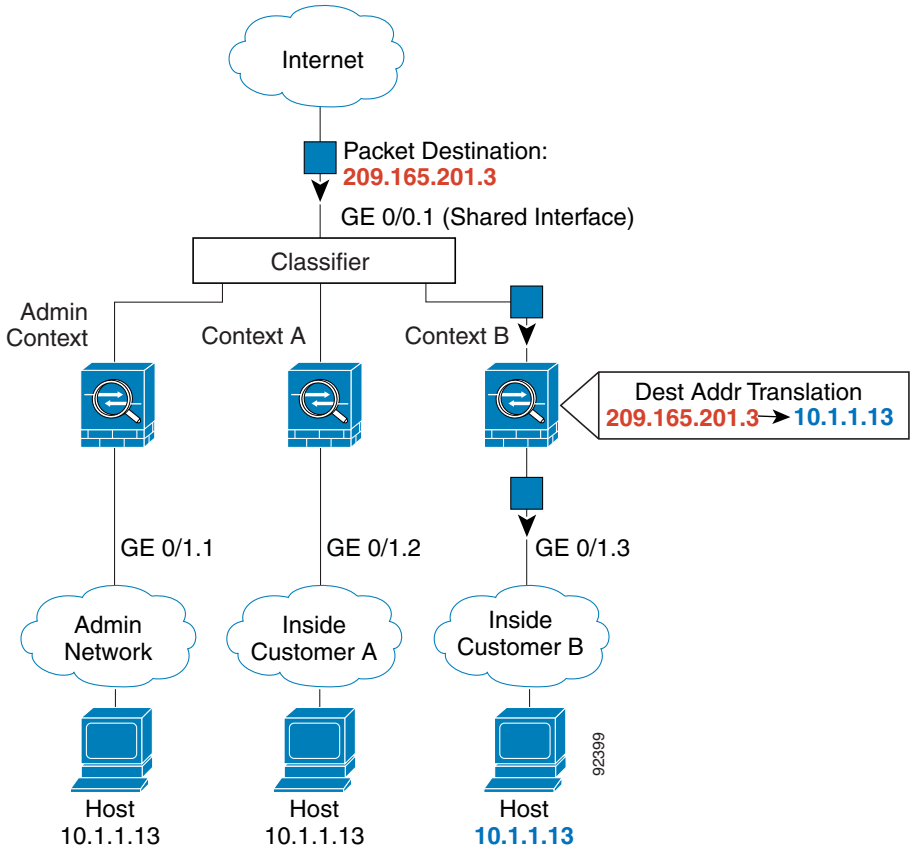

Note that all new incoming traffic must be classified, even from inside networks. Figure 7-3 shows a host on the Context B inside network accessing the Internet. The classifier assigns the packet to Context B because the ingress interface is Gigabit Ethernet 0/1.3, which is assigned to Context B.

Note If you share an *inside* interface and do not use unique MAC addresses, the classifier imposes some major restrictions. The classifier relies on the address translation configuration to classify the packet within a context, and you must translate the *destination* addresses of the traffic. Because you do not usually perform NAT on outside addresses, sending packets from inside to outside on a shared interface is not always possible; the outside network is large, (the Web, for example), and addresses are not predictable for an outside NAT configuration. If you share an inside interface, we suggest you use unique MAC addresses.

Figure 7-3 Incoming Traffic from Inside Networks

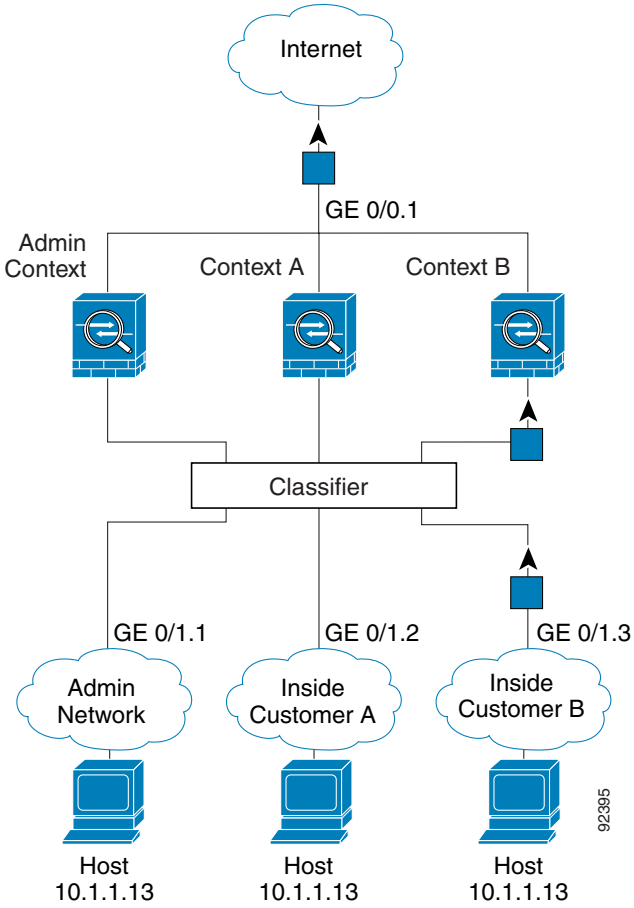

For transparent firewalls, you must use unique interfaces. Figure 7-4 shows a host on the Context B inside network accessing the Internet. The classifier assigns the packet to Context B because the ingress interface is Gigabit Ethernet 1/0.3, which is assigned to Context B.

Figure 7-4 Transparent Firewall Contexts

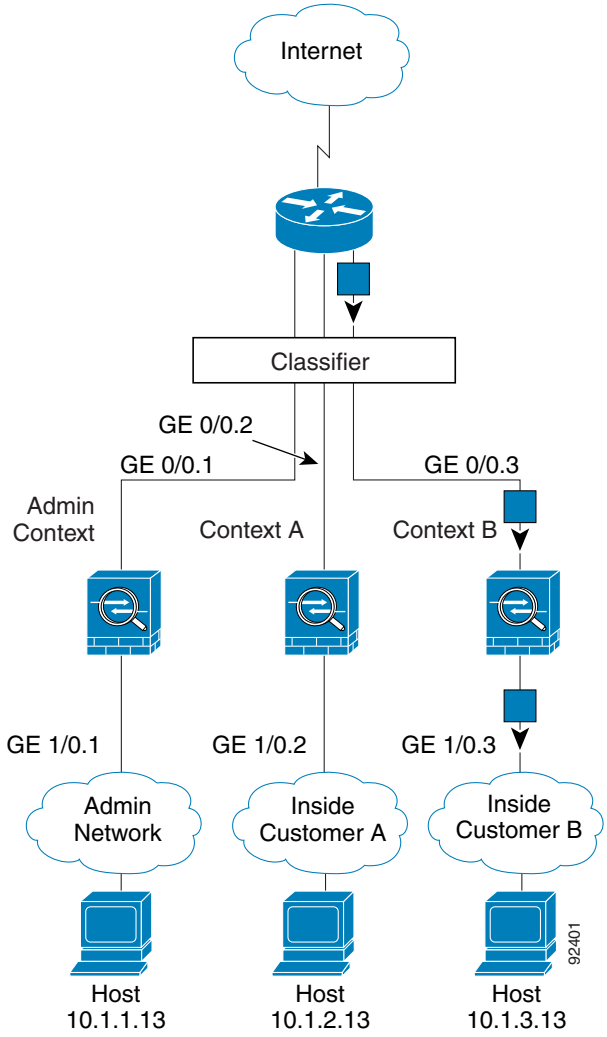

Cascading Security Contexts

Placing a context directly in front of another context is called cascading contexts; the outside interface of one context is the same interface as the inside interface of another context. You might want to cascade contexts if you want to simplify the configuration of some contexts by configuring shared parameters in the top context.

Note Cascading contexts requires that you configure unique MAC addresses for each context interface. Because of the limitations of classifying packets on shared interfaces without MAC addresses, we do not recommend using cascading contexts without unique MAC addresses.

Figure 7-5 shows a gateway context with two contexts behind the gateway.

Figure 7-5 Cascading Contexts

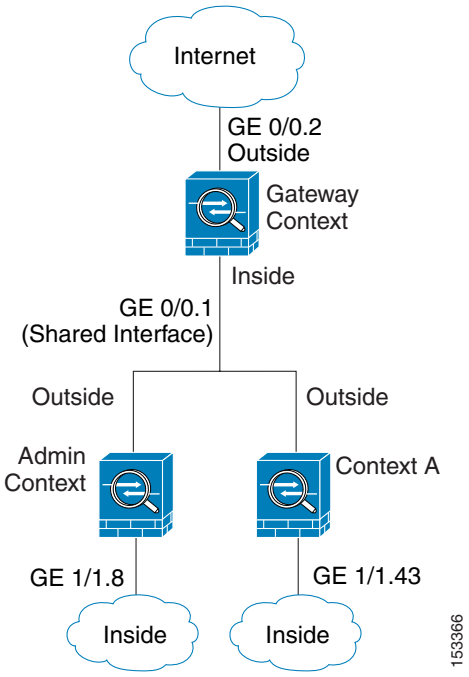

Management Access to Security Contexts

The security appliance provides system administrator access in multiple context mode as well as access for individual context administrators. The following sections describe logging in as a system administrator or as a a context administrator:

- System Administrator Access, page 7-8
- Context Administrator Access, page 7-9

System Administrator Access

You can access the security appliance as a system administrator in two ways:

- Access the security appliance console.

  From the console, you access the system execution space.

- Access the admin context using Telnet, SSH, or ASDM.

  See Chapter 11, "Configuring Device Access," to enable Telnet, SSH, and SDM access.

As the system administrator, you can access all contexts.

When you change to a context from admin or the system, your username changes to the default "enable_15" username. If you configured command authorization in that context, you need to either configure authorization privileges for the "enable_15" user, or you can log in as a different name for which you provide sufficient privileges in the command authorization configuration for the context. To log in with a username, enter the **login** command. For example, you log in to the admin context with the

username “admin.” The admin context does not have any command authorization configuration, but all other contexts include command authorization. For convenience, each context configuration includes a user “admin” with maximum privileges. When you change from the admin context to context A, your username is altered, so you must log in again as “admin” by entering the **login** command. When you change to context B, you must again enter the **login** command to log in as “admin.”

The system execution space does not support any AAA commands, but you can configure its own enable password, as well as usernames in the local database to provide individual logins.

Context Administrator Access

You can access a context using Telnet, SSH, or ASDM. If you log in to a non-admin context, you can only access the configuration for that context. You can provide individual logins to the context. See Chapter 11, “Configuring Device Access,” to enable Telnet, SSH, and SDM access and to configure management authentication.

Enabling or Disabling Multiple Context Mode at the CLI

Your security appliance might already be configured for multiple security contexts depending on how you ordered it from Cisco. If you are upgrading, however, you might need to convert from single mode to multiple mode by following the procedures in this section. ASDM does not support changing modes, so you need to change modes using the CLI.

This section includes the following topics:

Backing Up the Single Mode Configuration

When you convert from single mode to multiple mode, the security appliance converts the running configuration into two files. The original startup configuration is not saved, so if it differs from the running configuration, you should back it up before proceeding.

Enabling Multiple Context Mode

The context mode (single or multiple) is not stored in the configuration file, even though it does endure reboots. If you need to copy your configuration to another device, set the mode on the new device to match using the **mode** command.

When you convert from single mode to multiple mode, the security appliance converts the running configuration into two files: a new startup configuration that comprises the system configuration, and admin.cfg that comprises the admin context (in the root directory of the internal Flash memory). The original running configuration is saved as old_running.cfg (in the root directory of the internal Flash memory). The original startup configuration is not saved. The security appliance automatically adds an entry for the admin context to the system configuration with the name “admin.”

To enable multiple mode, enter the following command:

```
hostname(config)# mode multiple
```

You are prompted to reboot the security appliance.

Restoring Single Context Mode

If you convert from multiple mode to single mode, you might want to first copy a full startup configuration (if available) to the security appliance; the system configuration inherited from multiple mode is not a complete functioning configuration for a single mode device. Because the system configuration does not have any network interfaces as part of its configuration, you must access the security appliance from the console to perform the copy.

To copy the old running configuration to the startup configuration and to change the mode to single mode, perform the following steps in the system execution space:

Step 1 To copy the backup version of your original running configuration to the current startup configuration, enter the following command in the system execution space:

```
hostname(config)# copy flash:old_running.cfg startup-config
```

Step 2 To set the mode to single mode, enter the following command in the system execution space:

```
hostname(config)# mode single
```

The security appliance reboots.

Configuring Resource Classes

By default, all security contexts have unlimited access to the resources of the security appliance, except where maximum limits per context are enforced. However, if you find that one or more contexts use too many resources, and they cause other contexts to be denied connections, for example, then you can configure resource management to limit the use of resources per context.

This section includes the following topics:

Classes and Class Members Overview

The security appliance manages resources by assigning contexts to resource classes. Each context uses the resource limits set by the class. This section includes the following topics:

Resource Limits

When you create a class, the security appliance does not set aside a portion of the resources for each context assigned to the class; rather, the security appliance sets the maximum limit for a context. If you oversubscribe resources, or allow some resources to be unlimited, a few contexts can "use up" those resources, potentially affecting service to other contexts.

You can set the limit for individual resources, as a percentage (if there is a hard system limit) or as an absolute value.

You can oversubscribe the security appliance by assigning more than 100 percent of a resource across all contexts. For example, you can set the Bronze class to limit connections to 20 percent per context, and then assign 10 contexts to the class for a total of 200 percent. If contexts concurrently use more than the system limit, then each context gets less than the 20 percent you intended. (See Figure 7-6.)

Figure 7-6 Resource Oversubscription

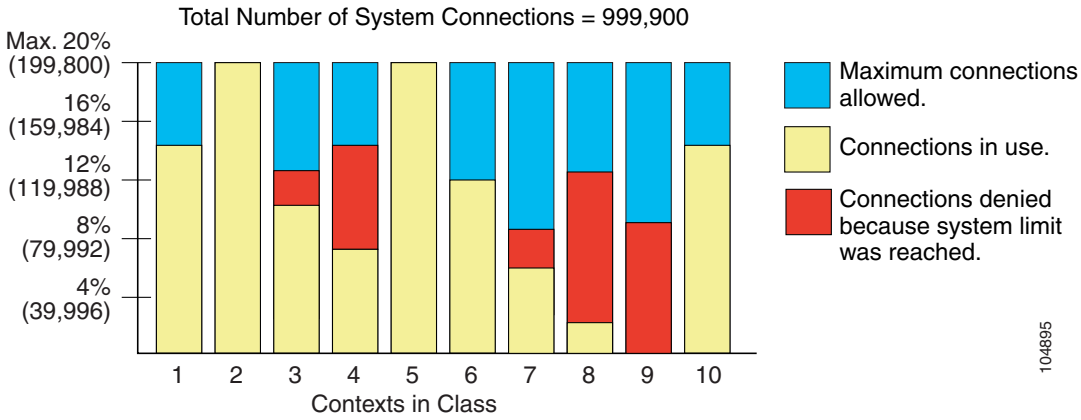

If you assign an absolute value to a resource across all contexts that exceeds the practical limit of the security appliance, then the performance of the security appliance might be impaired.

The security appliance lets you assign unlimited access to one or more resources in a class, instead of a percentage or absolute number. When a resource is unlimited, contexts can use as much of the resource as the system has available or that is practically available. For example, Context A, B, and C are in the Silver Class, which limits each class member to 1 percent of the connections, for a total of 3 percent; but the three contexts are currently only using 2 percent combined. Gold Class has unlimited access to connections. The contexts in the Gold Class can use more than the 97 percent of "unassigned" connections; they can also use the 1 percent of connections not currently in use by Context A, B, and C, even if that means that Context A, B, and C are unable to reach their 3 percent combined limit. (See Figure 7-7.) Setting unlimited access is similar to oversubscribing the security appliance, except that you have less control over how much you oversubscribe the system.

Figure 7-7 Unlimited Resources

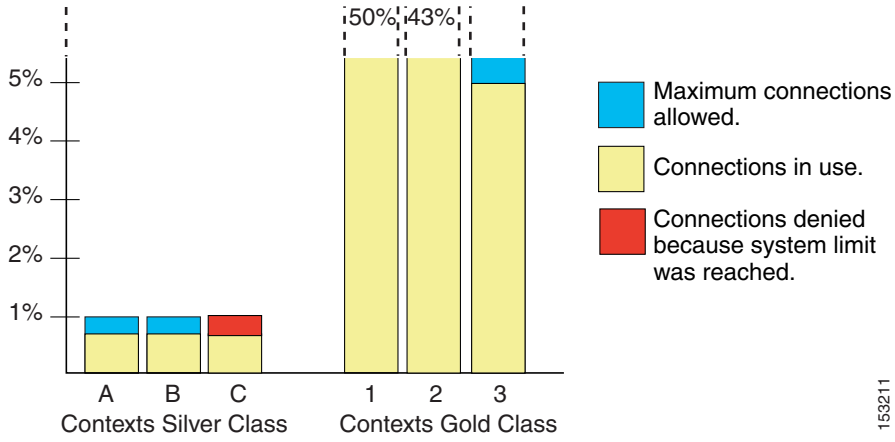

Default Class

All contexts belong to the default class if they are not assigned to another class; you do not have to actively assign a context to the default class.

If a context belongs to a class other than the default class, those class settings always override the default class settings. However, if the other class has any settings that are not defined, then the member context uses the default class for those limits. For example, if you create a class with a 2 percent limit for all concurrent connections, but no other limits, then all other limits are inherited from the default class. Conversely, if you create a class with a limit for all resources, the class uses no settings from the default class.

By default, the default class provides unlimited access to resources for all contexts, except for the following limits, which are by default set to the maximum allowed per context:

- Telnet sessions—5 sessions.
- SSH sessions—5 sessions.
- IPSec sessions—5 sessions.
- MAC addresses—65,535 entries.

Figure 7-8 shows the relationship between the default class and other classes. Contexts A and C belong to classes with some limits set; other limits are inherited from the default class. Context B inherits no limits from default because all limits are set in its class, the Gold class. Context D was not assigned to a class, and is by default a member of the default class.

Figure 7-8 Resource Classes

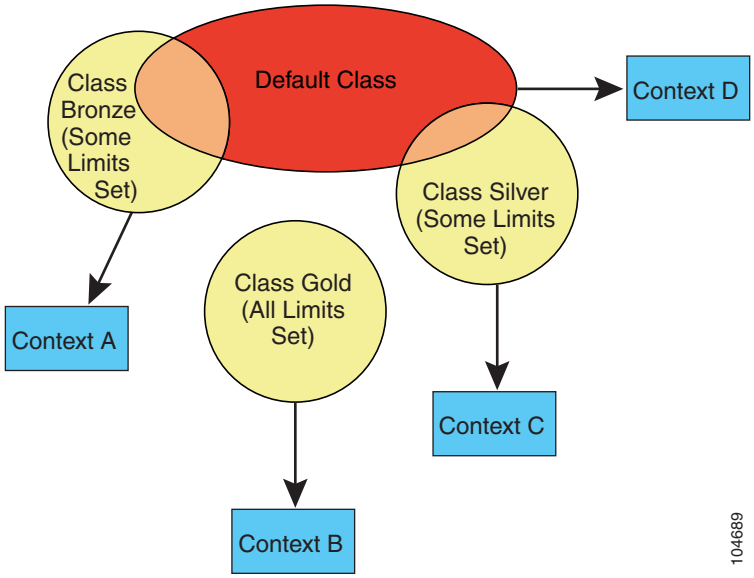

Class Members

To use the settings of a class, assign the context to the class when you define the context. All contexts belong to the default class if they are not assigned to another class; you do not have to actively assign a context to default. You can only assign a context to one resource class. The exception to this rule is that limits that are undefined in the member class are inherited from the default class; so in effect, a context could be a member of default plus another class.

Adding a Resource Class

This section describes the panes available for configuring resource classes, and includes the following topics:

- Resource Class, page 7-13
- Add/Edit Resource Class, page 7-14

Resource Class

System > Configuration > Resource Class

The Resource Class pane shows the configured classes and information about each class. It also lets you add, edit, or delete a class.

Fields

- Class—Shows the class name.
- All Resources—Shows the limit for all resources that you do not set individually, which can only be 0, which means unlimited.
- Connections—Shows the limit for TCP or UDP connections between any two hosts, including connections between one host and multiple other hosts.
- Hosts—Shows the limit for hosts that can connect through the security appliance.
- Xlates—Shows the limit for address translations.
- Telnet—Shows the limit for Telnet sessions, by default 5.
- SSH—Shows the limit for SSH sessions, by default 5.
- ASDM Sessions—Shows the limit for ASDM management sessions, by default 5. ASDM sessions use two HTTPS connections: one for monitoring that is always present, and one for making configuration changes that is present only when you make changes. For example, the system limit of 32 ASDM sessions represents a limit of 64 HTTPS sessions, divided between all contexts.
- MAC—Shows the limit for MAC addresses in the MAC address table in transparent firewall mode, by default 65535.
- Conns/sec—Shows the limit for connections per second.
- Fixups/sec—Shows the limit for application inspections per second.
- Syslogs/sec—Shows the limit for system log messages per second.
- Contexts—Shows the contexts assigned to this class.
- Add—Adds a class.
- Edit—Edits a class.
- Delete—Deletes a class. You cannot delete the default class. If you delete a class to which you assigned contexts, the contexts revert to using the default class.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

Add/Edit Resource Class

System > Configuration > Security Contexts > Add/Edit Resource Class

The Add/Edit Resource Class dialog box lets you add or edit a resource class.

Fields

- Resource Class—Sets the class name as a string up to 20 characters in length.

- Count Limited Resources—Sets the concurrent limits for resources. For resources that do not have a system limit, you cannot set the percentage; you can only set an absolute value. If you do not set a limit, the limit is inherited from the default class. If the default class does not set a limit, then the resource is unlimited, or the system limit if available.
  - Hosts—Sets the limit for concurrent hosts that can connect through the security appliance. Select the check box to enable this limit. If you set the limit to 0, it is unlimited.
  - Telnet—Sets the limit for concurrent Telnet sessions. Select the check box to enable this limit. You can set the limit as a percentage by entering any integer greater than 1 and selecting **Percent** from the list. You can assign more than 100 percent if you want to oversubscribe the device. Or you can set the limit as an absolute value by entering an integer between 1 and 5 and selecting **Absolute** from the list. The system has a maximum of 100 sessions divided between all contexts.
  - ASDM Sessions—Sets the limit for concurrent ASDM sessions. Select the check box to enable this limit. You can set the limit as a percentage by entering any integer greater than 1 and selecting **Percent** from the list. You can assign more than 100 percent if you want to oversubscribe the device. Or you can set the limit as an absolute value by entering an integer between 1 and 5 and selecting **Absolute** from the list. The system has a maximum of 80 sessions divided between all contexts. ASDM sessions use two HTTPS connections: one for monitoring that is always present, and one for making configuration changes that is present only when you make changes. For example, the system limit of 32 ASDM sessions represents a limit of 64 HTTPS sessions, divided between all contexts.
  - Connections—Sets the limit for concurrent TCP or UDP connections between any two hosts, including connections between one host and multiple other hosts. Select the check box to enable this limit. You can set the limit as a percentage by entering any integer greater than 1 and selecting **Percent** from the list. You can assign more than 100 percent if you want to oversubscribe the device. Or you can set the limit as an absolute value by entering an integer between 0 (system limit) and the system limit for your model, and selecting **Absolute** from the list. See the *Cisco ASDM Release Notes* for the connection limit for your model.
  - Xlates—Sets the limit for address translations. Select the check box to enable this limit. If you set the limit to 0, it is unlimited.
  - SSH—Sets the limit for SSH sessions. Select the check box to enable this limit. You can set the limit as a percentage by entering any integer greater than 1 and selecting **Percent** from the list. You can assign more than 100 percent if you want to oversubscribe the device. Or you can set the limit as an absolute value by entering an integer between 1 and 5 and selecting **Absolute** from the list. The system has a maximum of 100 sessions divided between all contexts.
  - MAC Entries—(Transparent mode only) Sets the limit for MAC address entries in the MAC address table. Select the check box to enable this limit. You can set the limit as a percentage by entering any integer greater than 1 and selecting **Percent** from the list. You can assign more than 100 percent if you want to oversubscribe the device. Or you can set the limit as an absolute value by entering an integer between 0 (system limit) and 65535 and selecting **Absolute** from the list.
- Rate Limited Resources—Sets the rate limit for resources. If you do not set a limit, the limit is inherited from the default class. If the default class does not set a limit, then it is unlimited by default.
  - Conns/sec—Sets the limit for connections per second. Select the check box to enable this limit. If you set the limit to 0, it is unlimited.
  - Syslogs/sec—Sets the limit for system log messages per second. Select the check box to enable this limit. If you set the limit to 0, it is unlimited.

  - Inspects/sec—Sets the limit for application inspections per second. Select the check box to enable this limit. If you set the limit to 0, it is unlimited.

- Show Actual Class Limits—(Non-default classes only) When you edit a class, this button shows the limits you set plus any inherited limits from the default class for limits you did not set.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

Configuring Security Contexts

This section describes how to add security contexts, and includes the following topics:

- Security Contexts, page 7-16
- Add/Edit Context, page 7-18
- Add/Edit Interface Allocation, page 7-18

Security Contexts

System > Configuration > Security Contexts

The Security Contexts pane shows the configured contexts and information about each context. It also lets you add, edit, or delete a context. For more information about multiple context mode, see the "Security Context Overview" section on page 7-1.

Prerequisites

Before you can configure contexts using ASDM, make sure the security appliance is in multiple context mode. If the ASDM toolbar includes Context and System tools, then the security appliance is in multiple mode. Also, the Home > Device Information > General tab shows the current context mode, either multiple or single. To change from single mode to multiple, access the security appliance CLI and enter the **mode multiple** command. See the "Enabling or Disabling Multiple Context Mode at the CLI" section on page 7-9 for more information.

Fields

- Context—Shows the context name.
- Interfaces—Shows the interfaces and subinterfaces assigned to the context. If you assigned an alias for the interface name to show in a context, then the aliased name is shown in parentheses. If you specified a range of subinterfaces, the range displays with a dash between the first and last subinterface numbers.
- Resource—Shows the resource class for each context.
- Config URL—Shows the context configuration location.

- Group—Shows the failover group to which this context belongs.
- Description—Shows a description of the context.
- Add—Adds a context.
- Edit—Edits a context.
- Delete—Deletes a context.
- Mac-Address auto—Automatically assigns private MAC addresses to each shared context interface.

  To allow contexts to share interfaces, we suggest that you assign unique MAC addresses to each context interface. The MAC address is used to classify packets within a context. If you share an interface, but do not have unique MAC addresses for the interface in each context, then the destination IP address is used to classify packets. The destination address is matched with the context NAT configuration, and this method has some limitations compared to the MAC address method. See the "How the Security Appliance Classifies Packets" section on page 7-2 for information about classifying packets.

  By default, the physical interface uses the burned-in MAC address, and all subinterfaces of a physical interface use the same burned-in MAC address.

  For use with failover, the security appliance generates both an active and standby MAC address for each interface. If the active unit fails over and the standby unit becomes active, the new active unit starts using the active MAC addresses to minimize network disruption.

  When you assign an interface to a context, the new MAC address is generated immediately. If you enable this option after you create context interfaces, then MAC addresses are generated for all interfaces immediately after you apply the option. If you disable this option, the MAC address for each interface reverts to the default MAC address. For example, subinterfaces of GigabitEthernet 0/1 revert to using the MAC address of GigabitEthernet 0/1.

  The MAC address is generated using the following format:

  Active unit MAC address: 12*_slot.port_subid.contextid*.

  Standby unit MAC address: 02*_slot.port_subid.contextid*.

  For platforms with no interface slots, the slot is always 0. The *port* is the interface port. The *subid* is an internal ID for the subinterface, which is not viewable. The *contextid* is an internal ID for the context. For example, the interface GigabitEthernet 0/1.200 in the context with the ID 1 has the following generated MAC addresses, where the internal ID for subinterface 200 is 31:

  Active: 1200.0131.0001

  Standby: 0200.0131.0001

  In the rare circumstance that the generated MAC address conflicts with another private MAC address in your network, you can manually set the MAC address for the interface within the context. See the "Configuring the Interfaces" section on page 4-2 to manually set the MAC address.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

Add/Edit Context

System > Configuration > Security Contexts > Add/Edit Context

The Add Context dialog box lets you add or edit a security context and define context parameters.

Fields

- Security Context—Sets the context name as a string up to 32 characters long. This name is case sensitive, so you can have two contexts named "customerA" and "CustomerA," for example. "System" or "Null" (in upper or lower case letters) are reserved names, and cannot be used.
- Interface Allocation—Shows the interfaces and subinterfaces assigned to this context.
  - Interface—Shows the interface IDs. If you specified a range of subinterfaces, the range displays with a dash between the first and last subinterface numbers.
  - Aliased Name—Shows the aliased name for this interface to be used in the context configuration instead of the interface ID.
  - Visible—Shows whether context users can see physical interface properties even if you set an aliased name.
  - Add—Adds an interface to the context.
  - Edit—Edits the interface properties.
  - Delete—Deletes an interface.
- Resource Assignment—Assigns the context to a resource class.
  - Resource Class—Select a class from the list.
  - Edit—Edits the selected resource class.
  - New—Adds a resource class.
- Config URL—Specifies the context configuration location, as a URL. Choose the file system type in the list, and then enter the server (for remote file systems), path, and filename in the field. For example, the combined URL for FTP has the following format:

  ftp://server.example.com/configs/admin.cfg
- Login—Sets the username and password for remote file systems.
- Failover Group—Sets the failover group for active/active failover.
- Description—Sets an optional description for the context.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

Add/Edit Interface Allocation

System > Configuration > Security Contexts > Add/Edit Context > Add/Edit Interface Allocation

The Add/Edit Interface Allocation dialog box lets you assign interfaces to a context and set interface parameters.

Fields

- Interfaces—Specifies the physical interface and subinterface IDs.
  - Physical Interface—Sets the physical interface to assign to the context. You can assign the main interface, in which case you leave the subinterface ID blank, or you can assign a subinterface or a range of subinterfaces associated with this interface. In transparent firewall mode, only interfaces that have not been allocated to other contexts are shown. If the main interface was already assigned to another context, then you must choose a subinterface.
  - Sub Interface Range (Optional)—Sets the subinterface ID or a range of subinterface IDs. To specify a single subinterface, choose the ID in the first list. To specify a range, choose the ending ID in the second list, if available. In transparent firewall mode, only subinterfaces that have not been allocated to other contexts are shown.
- Aliased Names—Sets an aliased name for this interface to be used in the context configuration instead of the interface ID.
  - Use Aliased Name in Context—Enables aliased names in the context.
  - Name—Sets the aliased name. An aliased name must start with a letter, end with a letter or digit, and have as interior characters only letters, digits, or an underscore. This field lets you specify a name that ends with a letter or underscore; to add an optional digit after the name, set the digit in the Range field.
  - Range—Sets the numeric suffix for the aliased name. If you have a range of subinterfaces, you can enter a range of digits to be appended to the name.
- Show Hardware Properties in Context—Enables context users to see physical interface properties even if you set an aliased name.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

Configuring Device Properties

This section contins the following topics:

Management IP

Configuration > Properties > Management IP

The Management IP window lets you set the management IP address for the security appliance or for a context in transparent firewall mode. A transparent firewall does not participate in IP routing. The only IP configuration required for the security appliance is to set the management IP address. The exception is that you can set the IP address for the Management 0/0 management-only interface, which does not pass through traffic. See the Configuring the Interfaces to set the IP address for Management 0/0.

This address is required because the security appliance uses this address as the source address for traffic originating on the security appliance, such as system messages or communications with AAA servers. You can also use this address for remote management access.

Fields

- Management IP Address—Sets the management IP address.
- Subnet Mask—Sets the subnet mask.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
—	•	•	•	—

Device Administration

Under **Device Administration**, you can set basic parameters for the security appliance. This section contains the following topics:

Banner

Configuration > Properties > Device Administration > Banner

The **Banner** panel lets you configure message of the day, login, and session banners.

To create a banner, enter text into the appropriate box. Spaces in the text are preserved, however, tabs can be entered in the ASDM interface but cannot be entered through the command line interface. The tokens $(domain) and $(hostname) are replaced with the host name and domain name of the security appliance.

Use the $(hostname) and $(domain) tokens to echo the hostname and domain name specified in a particular context. Use the $(system) token to echo a banner configured in the system space in a particular context.

Multiple lines in a banner are handled by entering a line of text for each line you wish to add. Each line is then appended to the end of the existing banner. If the text is empty, then a carriage return (CR) will be added to the banner. There is no limit on the length of a banner other than RAM and Flash memory limits. You can only use ASCII characters, including new line (the Enter key, which counts as two characters).

When accessing the security appliance through Telnet or SSH, the session closes if there is not enough system memory available to process the banner messages or if a TCP write error occurs when attempting to display the banner messages.

To replace a banner, change the contents of the appropriate box and click **Apply**. To clear a banner, clear the contents of the appropriate box and click **Apply**.

Although the banner command is not available in the System Context through the ASDM panel, it can be configured with **Tools > Command Line Interface**.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Boot Image/Configuration

Configuration > Properties > Device Administration > Boot Image/Configuration

Boot Image/Configuration lets you choose which image file the security appliance will boot from, as well as which configuration file it will use at startup.

You can specify up to four local binary image files for use as the startup image, and one image located on a TFTP server for the device to boot from. If you specify an image located on a TFTP server, it must be first in the list. In the event the device cannot reach the tftp server to load the image from, it will attempt to load the next image file in the list located in Flash.

If you do not specify any boot variable, the first valid image on internal flash will be chosen to boot the system.

Fields

Boot Configuration

- **Boot Order**—Displays the order in which binary image files will be used to boot.
- **Boot Image Location**—Displays the physical location and path of the boot file.
- **Boot Config File Path**—Displays the location of the configuration file.
- **Add**—Lets you add a flash or tftp boot image entry to be used in the boot process.
- **Edit**—Lets you edit a flash or tftp boot image entry.
- **Delete**—Deletes the selected flash or tftp boot image entry.
- **Move Up**—Moves the selected flash or tftp boot image entry up in the boot order.
- **Move Down**—Moves the selected flash or tftp boot image entry down in the boot order.
- **Browse Flash**—Lets you specify the location of a boot image or configuration file.

ASDM Image Configuration

- **ASDM Image File Path**—Displays the location of the configuration file the device will use at startup.
- **Browse Flash**—Lets you specify the location of a boot image or configuration file.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

Add Boot Image

Configuration > Properties > Device Administration > Boot Image/Configuration > Add Boot Image

To add a boot image entry to the boot order list, click **Add** on the **Boot Image/Configuration** panel.

You can select a Flash or TFTP image to add a boot image to the boot order list.

Either type the path of the image, or click **Browse Flash** to specify the image location. You must type the path of the image location if you are using TFTP.

Fields

- **Flash Image**—Select to add a boot image located in the flash file system.
  - **Path**—Specify the path of the boot image in the flash file system.
- **TFTP Image**—Select to add a boot image located on a TFTP server.
  - **[Path]**—Enter the path on the TFTP server of the boot image file, including the IP address of the server.
- **OK**—Accepts changes and returns to the previous panel.
- **Cancel**—Discards changes and returns to the previous panel.
- **Help**—Provides more information.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

Clock

Configuration > Properties > Device Administration > Clock

The **Clock** panel lets you manually set the date and time for the security appliance. The time displays in the status bar at the bottom of the main ASDM window.

In multiple context mode, set the time in the system configuration only.

To dynamically set the time using an NTP server, see the **NTP** panel; time derived from an NTP server overrides any time set manually in the **Clock** panel.

Fields

- **Time Zone**—Sets the time zone as GMT plus or minus the appropriate number of hours. If you select the Eastern Time, Central Time, Mountain Time, or Pacific Time zone, then the time adjusts automatically for daylight saving time, from 2:00 a.m. on the first Sunday in April to 2:00 a.m. on the last Sunday in October.

  Note Changing the time zone on the security appliance may drop the connection to intelligent SSMs.

- **Date**—Sets the date. Select the date and year from the lists, and then click the day on the calendar.
- **Time**—Sets the time on a 24-hour clock.
  - **hh**, **mm**, and **ss** boxes—Sets the hour, minutes, and seconds.
- **Update Display Time**—Updates the time shown at the bottom right corner of the ASDM window. The current time updates automatically every ten seconds.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

Console

Configuration > Properties > Device Administration > Console

The **Console** panel lets you specify a time period in minutes for the management console to remain active. When it reaches the time limit you specify here, the console automatically shuts down.

Type the time period in the **Console Timeout** text box. To specify unlimited, enter 0. The default value is 0.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

Device

Configuration > Properties > Device Administration > Device

The **Device** panel lets you set the hostname and domain name for the security appliance.

The hostname appears as the command line prompt, and if you establish sessions to multiple devices, the hostname helps you keep track of where you enter commands. The hostname is also used in system messages.

For multiple context mode, the hostname that you set in the system execution space appears in the command line prompt for all contexts. The hostname that you optionally set within a context does not appear in the command line, can be used for a banner.

The security appliance appends the domain name as a suffix to unqualified names. For example, if you set the domain name to "example.com," and specify a syslog server by the unqualified name of "jupiter," then the security appliance qualifies the name to "jupiter.example.com."

Fields

- *Platform* **Host Name**—Sets the hostname. The default hostname depends on your platform.
- **Domain Name**—Sets the domain name. The default domain name is default.domain.invalid.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

FTP Mode

Configuration > Properties > Device Administration > FTP Mode

The **FTP Mode** panel configures FTP mode as active or passive. The security appliance can use FTP to upload or download image files or configuration files to or from an FTP server. In passive FTP, the client initiates both the control connection and the data connection. The server, which is the recipient of the data connection in passive mode, responds with the port number to which it is listening for the specific connection.

Fields

- **Specify FTP mode as passive**—Configures FTP mode as active or passive.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

ICMP Rules

Configuration > Properties > Device Administration > ICMP Rules

The **ICMP Rules** panel provides a table that lists the ICMP rules, which specify the addresses of all the hosts or networks that are allowed or denied ICMP access to the security appliance. You can use this table to add or change the hosts or networks that are allowed or prevented from sending ICMP messages to the security appliance.

The ICMP rule list controls ICMP traffic that terminates on any security appliance interface. If no ICMP control list is configured, then the security appliance accepts all ICMP traffic that terminates at any interface, including the outside interface. However, by default, the security appliance does not respond to ICMP echo requests directed to a broadcast address.

Note Use the **Security Policy** panel to configure access rules for ICMP traffic that is routed *through* the security appliance for destinations on a protected interface.

It is recommended that permission is always granted for the ICMP unreachable message type (type 3). Denying ICMP unreachable messages disables ICMP Path MTU discovery, which can halt IPSec and PPTP traffic. See RFC 1195 and RFC 1435 for details about Path MTU Discovery.

If an ICMP control list is configured, then the security appliance uses a first match to the ICMP traffic followed by an implicit deny all. That is, if the first matched entry is a permit entry, the ICMP packet continues to be processed. If the first matched entry is a deny entry or an entry is not matched, the security appliance discards the ICMP packet and generates a syslog message. An exception is when an ICMP control list is not configured; in that case, a **permit** statement is assumed.

Fields

- **Interface**—Lists the interface on the security appliance from which ICMP access is allowed.
- **Action**—Displays whether ICMP messages are permitted or not allowed from the specified network or host.
- **IP Address**—Lists the IP address of the network or host that is allowed or denied access.
- **Mask**—Lists the network mask associated with the network or host that is allowed access.
- **ICMP Type**—Lists the type of ICMP message to which the rule applies. Table 8-1 lists the supported ICMP type values.
- **Add**—Displays the **Add ICMP Rule** dialog box for adding a new ICMP rule to the end of the table.
- **Insert Before**—Adds an ICMP rule before the currently selected rule.
- **Insert After**—Adds an ICMP rule after the currently selected rule.
- **Edit**—Displays the **Edit ICMP Rule** dialog box for editing the selected host or network.
- **Delete**—Deletes the selected host or network.

Table 8-1 ICMP Type Literals

ICMP Type	Literal
0	echo-reply
3	unreachable
4	source-quench
5	redirect

Table 8-1 ICMP Type Literals (continued)

ICMP Type	Literal
6	alternate-address
8	echo
9	router-advertisement
10	router-solicitation
11	time-exceeded
12	parameter-problem
13	timestamp-request
14	timestamp-reply
15	information-request
16	information-reply
17	mask-request
18	mask-reply
31	conversion-error
32	mobile-redirect

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit ICMP Rule

Configuration > Properties > Device Administration > ICMP Rules > Add/Edit ICMP Rule

The **Add/Edit ICM Rule** dialog box lets you add or modify an ICMP rule, which specifies the addresses of all the hosts or networks that are allowed or denied ICMP access to the security appliance.

Fields

- **ICMP Type**—Specifies the type of ICMP message to which the rule applies. Table 8-2 lists the supported ICMP type values.
- **Interface**—Identifies the interface on the security appliance from which ICMP access is allowed.
- **IP Address**—Specifies the IP address of the network or host that is allowed or denied access.
- **Any Address**—Applies the action to all addresses received on the specified interface.
- **Mask**—Specifies the network mask associated with the network or host that is allowed access.
- **Action**—Specifies whether ICMP messages are permitted or not from the specified network or host.

- **Permit**—Causes ICMP messages from the specified host or network and interface to be allowed.
- **Deny**—Causes ICMP messages from the specified host or network and interface to be dropped.

Table 8-2 ICMP Type Literals

ICMP Type	Literal
0	echo-reply
3	unreachable
4	source-quench
5	redirect
6	alternate-address
8	echo
9	router-advertisement
10	router-solicitation
11	time-exceeded
12	parameter-problem
13	timestamp-request
14	timestamp-reply
15	information-request
16	information-reply
17	mask-request
18	mask-reply
31	conversion-error
32	mobile-redirect

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Management Access

Configuration > Properties > Device Administration > Management Access

The **Management Access** panel lets you enable or disable management access on a high-security interface and thus lets you perform management functions on the security appliance. With management access enabled, you can run ASDM on an internal interface with a fixed IP address over an IPSec VPN

tunnel. Use this feature if VPN is configured on the security appliance and the external interface is using a dynamically assigned IP address. For example, this feature is helpful for accessing and managing the security appliance securely from home using the VPN client.

Fields

- **Management Access Interface**—Lets you specify the interface to use for managing the security appliance. **None** disables management access and is the default. To enable management access, select the interface with the highest security, which will be an inside interface. You can enable management access on only one interface at a time.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

NTP

Configuration > Properties > Device Administration > NTP

The **NTP** panel lets you define NTP servers to dynamically set the time on the security appliance. The time displays in the status bar at the bottom of the main ASDM window.

Time derived from an NTP server overrides any time set manually in the **Clock** panel.

NTP is used to implement a hierarchical system of servers that provide a precisely synchronized time among network systems. This kind of accuracy is required for time-sensitive operations, such as validating CRLs, which include a precise time stamp. You can configure multiple NTP servers. The security appliance chooses the server with the lowest stratum—a measure of how reliable the data is.

Fields

- **NTP Server List**—Shows defined NTP servers.
  - **IP Address**—Shows the NTP server IP address.
  - **Interface**—Specifies the outgoing interface for NTP packets, if configured. The system does not include any interfaces, so it uses the admin context interfaces. If the interface is blank, then the security appliance uses the default admin context interface according to the routing table.
  - **Preferred?**—Shows whether this NTP server is a preferred server, Yes or No. NTP uses an algorithm to determine which server is the most accurate and synchronizes to that one. If servers are of similar accuracy, then the preferred server is used. However, if a server is significantly more accurate than the preferred one, the security appliance uses the more accurate one. For example, the security appliance uses a more accurate server over a less accurate server that is preferred.
  - **Key Number**—Shows the authentication key ID number.
  - **Trusted Key?**—Shows if the key is a trusted key, Yes or No. The key must be trusted for authentication to work.
- **Enable NTP Authentication**—Enables authentication for all servers.

- **Add**—Adds an NTP server.
- **Edit**—Edits an NTP server.
- **Delete**—Deletes and NTP server.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

Add/Edit NTP Server Configuration

Configuration > Properties > Device Administration > NTP > Add/Edit NTP Server Configuration

The **Add/Edit NTP Server Configuration** dialog box lets you add or edit an NTP server.

Fields

- **IP Address**—Sets the NTP server IP address.
- **Preferred**—Sets this server as a preferred server. NTP uses an algorithm to determine which server is the most accurate and synchronizes to that one. If servers are of similar accuracy, then the preferred server is used. However, if a server is significantly more accurate than the preferred one, the security appliance uses the more accurate one. For example, the security appliance uses a more accurate server over a less accurate server that is preferred.
- **Interface**—Sets the outgoing interface for NTP packets, if you want to override the default interface according to the routing table. The system does not include any interfaces, so it uses the admin context interfaces. If you intend to change the admin context (thus changing the available interfaces), you should choose **None** (the default interface) for stability.
- **Authentication Key**—Sets the authentication key attributes if you want to use MD5 authentication for communicating with the NTP server.
  - **Key Number**—Sets the key ID for this authentication key. The NTP server packets must also use this key ID. If you previously configured a key ID for another server, you can select it in the list; otherwise, type a number between 1 and 4294967295.
  - **Trusted**—Sets this key as a trusted key. You must select this box for authentication to work.
  - **Key Value**—Sets the authentication key as a string up to 32 characters in length.
  - **Reenter Key Value**—Validates the key by ensuring that you enter the key correctly two times.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

Password

Configuration > Properties > Device Administration > Password

The **Password** panel lets you set the login password and the enable password.

The login password lets you access EXEC mode if you connect to the security appliance using a Telnet or SSH session. (If you configure user authentication for Telnet or SSH access, then each user has their own password, and this login password is not used; see the **AAA Access** panel.)

The enable password lets you access privileged EXEC mode after you log in. Also, this password is used to access ASDM as the default user, which is blank. The default user shows as "enable_15" in the **User Accounts** panel. (If you configure user authentication for enable access, then each user has their own password, and this enable password is not used; see the **AAA Access** panel. In addition, you can configure authentication for HTTP/ASDM access.)

Fields

- **Enable Password**—Sets the enable password. By default, it is blank.
  - **Change the privileged mode password**—Lets you change the enable password.
  - **Old Password**—Enter the old password.
  - **New Password**—Enter the new password.
  - **Confirm New Password**—Confirm the new password.
- **Telnet Password**—Sets the login password. By default, it is "cisco." Although this group box is called Telnet Password, this password applies to Telnet and SSH access.
  - **Change the password to access the** *platform* **console**—Lets you change the login password.
  - **Old Password**—Enter the old password.
  - **New Password**—Enter the new password.
  - **Confirm New Password**—Confirm the new password.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Secure Copy

Configuration > Properties > Device Administration > Secure Copy

The **Secure Copy** panel lets you enable the secure copy server on the security appliance. Only clients that are allowed to access the security appliance using SSH can establish a secure copy connection.

Limitations

This implementation of the secure copy server has the following limitations:

- The server can accept and terminate connections for secure copy, but cannot initiate them.
- The server does not have directory support. The lack of directory support limits remote client access to the security appliance internal files.
- The server does not support banners.
- The server does not support wildcards.
- The security appliance license must have the VPN-3DES-AES feature to support SSH version 2 connections.

Fields

- **Enable Secure Copy Server**—Select this check box to enable the secure copy server on the security appliance.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

SMTP

Configuration > Properties > Device Administration > SMTP

The **SMTP** panel lets you enable or disable the SMTP client for notification by email that a significant event has transpired. Here you can add an IP address of an SMTP server and optionally, the IP address of a backup server. ASDM does not check to make sure the IP address is valid, so it is important to type the address correctly.

You can configure what email addresses will receive alerts in **Configuration > Properties > Logging > Email Setup**.

Fields

- **Remote SMTP Server**—Lets you configure the primary and secondary SMTP servers.
- **Primary Server IP Address**—Enter the IP address of the SMTP server.
- **Secondary Server IP Address (Optional)**—Optionally, you can enter the IP address of a secondary SMTP server.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

SNMP

Configuration > Properties > Device Administration > SNMP

The **SNMP** panel lets you configure the security appliance for monitoring by Simple Network Management Protocol (SNMP) management stations.

SNMP defines a standard way for network management stations running on PCs or workstations to monitor the health and status of many types of devices, including switches, routers, and the security appliance.

SNMP Terminology

- **Management station**—Network management stations running on PCs or workstations, use the SNMP protocol to administer standardized databases residing on the device being managed. Management stations can also receive messages about events, such as hardware failures, which require attention.
- **Agent**—In the context of SNMP, the management station is a *client* and an SNMP agent running on the security appliance is a *server*.
- **OID**—The SNMP standard assigns a system object ID (OID) so that a management station can uniquely identify network devices with SNMP agents and indicate to users the source of information monitored and displayed.
- **MIB**—The agent maintains standardized data structures called Management Information Databases, or MIBs which are compiled into management stations. MIBs collect information, such as packet, connection, and error counters, buffer usage, and failover status. MIBs are defined for specific products, in addition to MIBs for the common protocols and hardware standards used by most network devices. SNMP management stations can browse MIBs or request only specific fields. In some applications, MIB data can be modified for administrative purposes.
- **Trap**—The agent also monitors alarm conditions. When an alarm condition defined in a trap occurs, such as a link up, link down, or syslog event, the agent sends notification, also known as SNMP trap, to the designated management station immediately.

SNMP

For Cisco MIB files and OIDs, refer to: http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml. OIDs may be downloaded at this URL: ftp://ftp.cisco.com/pub/mibs/oid/oid.tar.gz.

MIB Support

The security appliance provides the following SNMP MIB support:

Note The security appliance does not support browsing of the Cisco syslog MIB.

- You can browse the System and Interface groups of MIB-II. Browsing an MIB is different from sending traps. Browsing means doing an snmpget or snmpwalk of the MIB tree from the management station to determine values.
- The Cisco MIB and Cisco Memory Pool MIB are available.
- The security appliance does not support the following in the Cisco MIB:
- cfwSecurityNotification NOTIFICATION-TYPE
- cfwContentInspectNotification NOTIFICATION-TYPE
- cfwConnNotification NOTIFICATION-TYPE
- cfwAccessNotification NOTIFICATION-TYPE
- cfwAuthNotification NOTIFICATION-TYPE
- cfwGenericNotification NOTIFICATION-TYPE

SNMP CPU Utilization

The security appliance supports monitoring CPU utilization through SNMP. This feature allows network administrators to monitor security appliance CPU usage using SNMP management software, such as HP OpenView, for capacity planning.

This functionality is implemented through support for the cpmCPUTotalTable of the Cisco Process MIB (CISCO-PROCESS-MIB.my). The other two tables in the MIB, cpmProcessTable and cpmProcessExtTable, are not supported in this release.

Each row of the cpmCPUTotalTable includes the index of each CPU and the following objects:

MIB object name	Description
cpmCPUTotalPhysicalIndex	The value of this object will be zero because the entPhysicalTable of Entity MIB is not supported on the security appliance SNMP agent.
cpmCPUTotalIndex	The value of this object will be zero because the entPhysicalTable of Entity MIB is not supported on the security appliance SNMP agent.
cpmCPUTotal5sec	Overall CPU busy percentage in the last five-second period.
cpmCPUTotal1min	Overall CPU busy percentage in the last one-minute period.
cpmCPUTotal5min	Overall CPU busy percentage in the last five-minute period.

Note Because all current security appliance hardware platforms support a single CPU, the security appliance returns only one row from cpmCPUTotalTable and the index is always 1.

The values of the last three elements are the same as the output from the show cpu usage command.

The security appliance does not support the following new MIB objects in the cpmCPUTotalTable:

- cpmCPUTotal5secRev
- cpmCPUTotal1minRev
- cpmCPUTotal5minRev

Fields

- **Community string (default)**—Enter the password used by the SNMP management station when sending requests to the security appliance. The SNMP community string is a shared secret among the SNMP management stations and the network nodes being managed. The security appliance uses the password to determine if the incoming SNMP request is valid. The password is a case-sensitive value up to 32 characters in length. Spaces are not permitted. The default is "public." SNMPv2c allows separate community strings to be set for each management station. If no community string is configured for any management station, the value set here will be used by default.
- **Contact**—Enter the name of the security appliance system administrator. The text is case-sensitive and can be up to 127 characters. Spaces are accepted, but multiple spaces are shortened to a single space.
- **Security Appliance Location**—Specify the security appliance location. The text is case-sensitive and can be up to 127 characters. Spaces are accepted, but multiple spaces are shortened to a single space.
- **Listening Port**—Specify the port on which SNMP traffic is sent. The default is 161.
- **Configure Traps**—Lets you configure the events to notify through SNMP traps.
- **SNMP Management Station** box:
  - **Interface**—Displays the security appliance interface name where the SNMP management station resides.
  - **IP Address**—Displays the IP address of an SNMP management station to which the security appliance sends trap events and receive requests or polls.
  - **Community string**—If no community string is specified for a management station, the value set in **Community String (default)** field will be used.
  - **SNMP Version**—Displays the version of SNMP set on the management station.
  - **Poll/Trap**—Displays the method for communicating with this management station, poll only, trap only, or both trap and poll. Polling means that the security appliance waits for a periodic request from the management station. The trap setting sends syslog events when they occur.
  - **UDP Port**—SNMP host UDP port. The default is port 162.
- **Add**—Opens **Add SNMP Host Access Entry** with these fields:
- **Interface Name**—Select the interface on which the management station resides.
- **IP Address**—Specify the IP address of the management station.
- **Server Poll/Trap Specification**—Select **Poll**, **Trap**, or both.
- **UDP Port**—UDP port for the SNMP host. This field allows you to override the default value of 162 for the SNMP host UDP port.
- **Help**—Provides more information.
- **Edit**—Opens the **Edit SNMP Host Access Entry** dialog box with the same fields as Add.
- **Delete**—Deletes the selected item.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit SNMP Host Access Entry

Configuration > Properties > Device Administration > SNMP > Add/Edit SNMP Host Access Entry

Adding SNMP Management Stations

To add SNMP management stations, perform the following steps:

1. Click **Add** to open the **SNMP Host Access Entry** dialog box.
2. From **Interface Name**, select the interface on which the SNMP management station resides.
3. Enter the IP address of that management station in **IP Address**.
4. Enter the UDP port for the SNMP host. The default is 162.
5. Enter the Community String password for the SNMP host. If no community string is specified for a management station, the value set in **Community String (default)** field in the SNMP Configuration screen will be used.
6. Click to select **Poll, Trap**, or both.
7. To return to the previous panel click:
   - **OK**—Accepts changes and returns to the previous panel
   - **Cancel**—Discards changes and returns to the previous panel
   - **Help**—Provides more information

Editing SNMP Management Stations

To edit SNMP management stations, perform the following steps:

1. Select a list item from the SNMP management station table on the **SNMP** panel.
2. Click **Edit** to open **Edit SNMP Host Access Entry**.
3. From **Interface Name**, select the interface on which the SNMP management station resides.
4. Enter the IP address of that management station in **IP Address**.
5. Enter the Community String password for the SNMP host. If no community string is specified for a management station, the value set in **Community String (default)** field in the SNMP Configuration screen will be used.
6. Enter the UDP port for the SNMP host. The default is 162.
7. Click to select **Poll**, **Trap**, or both.
8. Select SNMP version.
9. To return to the previous panel click:
   - **OK**—Accepts changes and returns to the previous panel
   - **Cancel**—Discards changes and returns to the previous panel

- **Help**—Provides more information

Deleting SNMP Management Stations

To delete an SNMP management station from the table, perform the following steps:

1. Select an item from the SNMP management station table on the **SNMP** panel.
2. Click **Delete**.

Fields

- **Interface name**—Select the interface where the SNMP host resides.
- **IP Address**—Enter the IP address of the SNMP host.
- **UDP Port**—Enter the UDP port on which to send SNMP updates. The default is 162.
- **Community String**—Enter the community string for the SNMP server.
- **SNMP Version**—Select the SNMP version.
- **Server Port/Trap Specification**
  - **Poll**—Select to send poll information. Polling means that the security appliance waits for a periodic request from the management station.
  - **Trap**—Select to send trap information. The trap setting sends syslog events when they occur.
- **OK**—Accepts changes and returns to the previous panel
- **Cancel**—Discards changes and returns to the previous panel
- **Help**—Provides more information

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

SNMP Trap Configuration

Configuration > Properties > Device Administration > SNMP > SNMP Trap Configuration

Traps

Traps are different than browsing; they are unsolicited “comments” from the managed device to the management station for certain events, such as link up, link down, and syslog event generated.

An SNMP object ID (OID) for the security appliance displays in SNMP event traps sent from the security appliance. The security appliance provides system OID in SNMP event traps & SNMP mib-2.system.sysObjectID.

The SNMP service running on the security appliance performs two different functions:

- Replies to SNMP requests from management stations (also known as SNMP clients).
- Sends traps (event notifications) to management stations or other devices that are registered to receive them from the security appliance.

The security appliance supports 3 types of traps:

- firewall
- generic
- syslog

Configure Traps

Opens **SNMP Trap Configuration** with the following fields:

- **Standard SNMP Traps**—Select standard traps to send:
  - **Authentication**—Enables authentication standard trap.
  - **Cold Start**—Enables cold start standard trap.
  - **Link Up**—Enables link up standard trap.
  - **Link Down**—Enables link down standard trap.
- **Entity MIB Notifications**
  - **FRU Insert**—Enables a trap notification when a Field Replaceable Unit (FRU) has been inserted.
  - **FRU Remove**—Enables a trap notification when a Field Replaceable Unit (FRU) has been removed.
  - **Configuration Change**—Enables a trap notification when there has been a hardware change.
- **IPSec Traps**—Enables IPSec traps.
  - **Start**—Enables a trap when IPSec starts.
  - **Stop**—Enables a trap when IPSec stops.
- **Remote Access Traps**—Enables remote access traps.
  - **Session threshold exceeded**—Enables a trap when the number of remote access session attempts exceeds the threshold configured.
- **Enable Syslog traps**—Enables sending of syslog messages to SNMP management station.
- **OK**—Accepts changes and returns to the previous panel.
- **Cancel**—Discards changes and returns to the previous panel.
- **Help**—Provides more information.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

TFTP Server

Configuration > Properties > Device Administration > TFTP

The **TFTP Server** panel lets you configure the security appliance to save its configuration to a file server using TFTP.

Note This panel does not write the file to the server. Configure the security appliance for using a TFTP server in this panel, then click **File > Save Running Configuration to TFTP Server**.

TFTP Servers and the security appliance

TFTP is a simple client/server file transfer protocol described in RFC783 and RFC1350 Rev. 2. This panel lets you configure the security appliance as a TFTP *client* so that it can transfer a copy of its running configuration file to a TFTP *server* using **File > Save Running Configuration to TFTP Server** or **Tools > Command Line Interface**. In this way, you can back up and propagate configuration files to multiple security appliances.

This panel uses the **configure net** command to specify the IP address of the TFTP server, and the **tftp-server** command to specify the interface and the path/filename on the server where the running configuration file will be written. Once this information is applied to the running configuration, ASDM **File > Save Running Configuration** to TFTP Server uses the **copy** command to execute the file transfer.

The security appliance supports only one TFTP server. The full path to the TFTP server is specified in **Configuration > Properties > Administration > TFTP Server**. Once configured here, you can use a colon (:) to specify the IP address in the CLI **configure net** and **copy** commands. However, any other authentication or configuration of intermediate devices necessary for communication from the security appliance to the TFTP server is done apart from this function.

The **show tftp-server** command lists the **tftp-server** command statements in the current configuration. The **no tftp server** command disables access to the server.

Fields

The **TFTP** panel provides the following fields:

- **Enable**—Click to select and enable these TFTP server settings in the configuration.
- **Interface Name**—Select the name of the security appliance interface which will use these TFTP server settings.
- **IP Address**—Enter the IP address of the TFTP server.
- **Path**—Type in the TFTP server path, beginning with "/" (forward slash) and ending in the file name, to which the running configuration file will be written.

  Example TFTP server path: **/tftpboot/**security appliance**/config3**

  Note The path must begin with a forward slash (/).

For More Information

For more information about TFTP, refer to the security appliance Technical Documentation for your version of software.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

User Accounts

Configuration > Properties > Device Administration > User Accounts

The **User Accounts** panel lets you manage the local user database. The local database is used for the following features:

- ASDM per-user access

  By default, you can log into ASDM with a blank username and the enable password (see **Password**). However, if you enter a username and password at the login screen (instead of leaving the username blank), ASDM checks the local database for a match.

Note Although you can configure HTTP authentication using the local database (see the **Authentication Tab**), that functionality is always enabled by default. You should only configure HTTP authentication if you want to use a RADIUS or TACACS+ server for authentication.

- Console authentication (see the **Authentication Tab**)
- Telnet and SSH authentication (see the **Authentication Tab**)
- **enable** command authentication (see the **Authentication Tab**)

  This setting is for CLI-access only and does not affect the ASDM login.

- Command authorization (see the **Authorization Tab**)

  If you enable command authorization using the local database, then the security appliance refers to the user privilege level to determine what commands are available. Otherwise, the privilege level is not generally used. By default, all commands are either privilege level 0 or level 15. ASDM allows you to enable three predefined privilege levels, with commands assigned to level 15 (Admin), level 5 (Read Only), and level 3 (Monitor Only). If you use the predefined levels, then assign users to one of these three privilege levels.

Note If you add users to the local database who can gain access to the CLI and whom you do not want to enter privileged mode, you should enable command authorization. Without command authorization, users can access privileged mode (and all commands) at the CLI using their own password if their privilege level is 2 or greater (2 is the default). Alternatively, you can use RADIUS or TACACS+ authentication for console access so the user will not be able to use the login command, or you can set all local users to level 1 so you can control who can use the system enable password to access privileged mode.

- Network access authentication
- VPN client authentication

You cannot use the local database for network access authorization.

For multiple context mode, you can configure usernames in the system execution space to provide individual logins at the CLI using the **login** command; however, you cannot configure any **aaa** commands that use the local database in the system execution space.

Note VPN functions are not supported in multimode.

To configure the enable password from this panel (instead of in **Password**), change the password for the enable_15 user. The enable_15 user is always present in this panel, and represents the default username. This method of configuring the enable password is the only method available in ASDM for the system configuration. If you configured other enable level passwords at the CLI (**enable password 10**, for example), then those users are listed as enable_10, etc.

Fields

- **User Name**—Specifies the user name to which these parameters apply.
- **Privilege (Level)**—Specifies the privilege level assigned to that user. The privilege level is used with local command authorization. See the **Authorization Tab** for more information.
- **VPN Group Policy**—Specifies the name of the VPN group policy for this user. Not available in multimode.
- **VPN Group Lock**—Specifies what, if any, group lock policy is in effect for this user. Not available in multimode.
- **Add**—Displays the Add User Account dialog box.
- **Edit**—Displays the Edit User Account dialog box.
- **Delete**—Removes the selected row from the table. There is no confirmation or undo.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit User Account > Identity Tab

Configuration > Properties > Device Administration > User Accounts > Add/Edit User Account > Identity Tab

Use this tab to specify parameters that identify the user account you want to add or change. The changes appear in the User Accounts table as soon as you click OK.

Fields

- **Username**—Specifies the username for this account.
- **Password**—Specifies the unique password for this user. The minimum password length is 4 characters. The maximum is 32 characters. Entries are case-sensitive. The field displays only asterisks.

Note To protect security, we recommend a password length of at least 8 characters.

- **Confirm Password**—Asks you to re-enter the user password to verify it. The field displays only asterisks.
- **Privilege Level**—Selects the privilege level for this user to use with local command authorization. The range is 0 (lowest) to 15 (highest). See the **Authorization Tab** for more information.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit User Account > VPN Policy Tab

Configuration > Properties > Device Administration > User Accounts > Add/Edit User Account > VPN Policy Tab

Use this tab to specify VPN policies for this user. Check an Inherit check box to let the corresponding setting take its value from the group policy.

Fields

- **Group Policy**—Lists the available group policies.
- **Tunneling Protocols**—Specifies what tunneling protocols that this user can use, or whether to inherit the value from the group policy. Check the desired **Tunneling Protocols** check boxes to select the VPN tunneling protocols that this user can use. Users can use only the selected protocols. The choices are as follows:

  IPSec—IP Security Protocol. IPSec provides the most complete architecture for VPN tunnels, and it is perceived as the most secure protocol. Both LAN-to-LAN (peer-to-peer) connections and client-to-LAN connections can use IPSec.

  WebVPN—VPN via SSL/TLS. Uses a web browser to establish a secure remote-access tunnel to a VPN Concentrator; requires neither a software nor hardware client. WebVPN can provide easy access to a broad range of enterprise resources, including corporate websites, web-enabled applications, NT/AD file share (web-enabled), e-mail, and other TCP-based applications from almost any computer that can reach HTTPS Internet sites.

  L2TP over IPSec—Allows remote users with VPN clients provided with several common PC and mobile PC operating systems to establish secure connections over the public IP network to the security appliance and private corporate networks.

  Note If no protocol is selected, an error message appears.

- **Filter**—Specifies what filter to use, or whether to inherit the value from the group policy. Filters consist of rules that determine whether to allow or reject tunneled data packets coming through the security appliance, based on criteria such as source address, destination address, and protocol. To configure filters and rules, see the Configuration > VPN > VPN General > Group Policy panel.
- **Manage**—Displays the ACL Manager panel, on which you can add, edit, and delete Access Control Lists (ACLs) and Extended Access Control Lists (ACEs).
- **Tunnel Group Lock**—Specifies whether to inherit the tunnel group lock or to use the selected tunnel group lock, if any. Selecting a specific lock restricts users to remote access through this group only. Tunnel Group Lock restricts users by checking if the group configured in the VPN client is the same as the user's assigned group. If it is not, the security appliance prevents the user from connecting. If the Inherit check box is not selected, the default value is --None--.
- **Store Password on Client System**—Specifies whether to inherit this setting from the group. Deselecting the Inherit check box activates the Yes and No radio buttons. Selecting Yes stores the login password on the client system (potentially a less-secure option). Selecting No (the default) requires the user to enter the password with each connection. For maximum security, we recommend that you *not do allow* password storage. This parameter has no bearing on interactive hardware client authentication or individual user authentication for a VPN 3002.
- **Connection Settings**—Specifies the connection settings parameters.
  - **Access Hours**—If the Inherit check box is not selected,you can select the name of an existing access hours policy, if any, applied to this user or create a new access hours policy. The default value is Inherit, or, if the Inherit check box is not selected, the default value is --Unrestricted--.
  - **New**—Opens the Add Time Range dialog box, on which you can specify a new set of access hours.
  - **Simultaneous Logins**—If the Inherit check box is not selected, this parameter specifies the maximum number of simultaneous logins allowed for this user. The default value is 3. The minimum value is 0, which disables login and prevents user access.

    Note While there is no maximum limit, allowing several simultaneous connections could compromise security and affect performance.

  - **Maximum Connect Time**—If the Inherit check box is not selected, this parameter specifies the maximum user connection time in minutes. At the end of this time, the system terminates the connection. The minimum is 1 minute, and the maximum is 2147483647 minutes (over 4000 years). To allow unlimited connection time, select the Unlimited check box (the default).
  - **Idle Timeout**—If the Inherit check box is not selected, this parameter specifies this user's idle timeout period in minutes. If there is no communication activity on the user's connection in this period, the system terminates the connection. The minimum time is 1 minute, and the maximum time is 10080 minutes. This value does not apply to WebVPN users.
- **Dedicated IP Address (Optional)**—
  - **IP Address** box—Specifies the optional Dedicated IP address.
  - **Subnet Mask** list—Specifies the subnet mask for the Dedicated IP address.

Check the **Group Lock** check box to restrict users to remote access through this group only. Group Lock restricts users by checking if the group configured in the VPN client is the same as the user's assigned group. If it is not, the VPN Concentrator prevents the user from connecting.

If this box is unchecked (the default), the system authenticates a user without regard to the user's assigned group.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit User Account > WebVPN Tab

Configuration > Properties > Device Administration > User Accounts > Add/Edit User Account > WebVPN Tab

The **Add** or **Edit User Account** panel, **WebVPN** tab, displays six tabs that let you configure WebVPN attributes for users.

Fields

- **Inherit**—Indicates that the corresponding setting takes its value from the default group policy, rather than from the explicit specifications that follow.
- **Functions**—Configures the features available to WebVPN users.
  - **Enable URL entry**—Places the URL entry box on the home page. If this feature is enabled, users can enter web addresses in the URL entry box, and use WebVPN to access those websites.

    Using WebVPN does not ensure that communication with every site is secure. WebVPN ensures the security of data transmission between the remote user's PC or workstation and the security appliance on the corporate network. If a user then accesses a non-HTTPS web resource (located on the Internet or on the internal network), the communication from the corporate security appliance to the destination web server is not secured.

    In a WebVPN connection, the security appliance acts as a proxy between the end user's web browser and target web servers. When a WebVPN user connects to an SSL-enabled web server, the security appliance establishes a secure connection and validates the server's SSL certificate. The end user's browser never receives the presented certificate, so therefore cannot examine and validate the certificate. The current implementation of WebVPN does not permit communication with sites that present expired certificates. Neither does the security appliance perform trusted CA certificate validation. Therefore, WebVPN users cannot analyze the certificate an SSL-enabled web-server presents before communicating with it.

    To limit Internet access for WebVPN users, deselect the Enable URL Entry field. This prevents WebVPN users from surfing the Web during a WebVPN connection.
  - **Enable file server access**—Enables Windows file access (SMB/CIFS files only) through HTTPS. When this box is checked, users can access Windows files on the network. If you enable only this parameter for WebVPN file sharing, users can access only servers that you configure in Servers and URLs group box. To let users access servers directly or to browse servers on the network, see the Enable file server entry and Enable file server browsing parameters.

    Users can download, edit, delete, rename, and move files. They can also add files and folders.

    Shares must also be configured for user access on the applicable Windows servers. Users might have to be authenticated before accessing files, depending on network requirements.

File access, server/domain access, and browsing require that you configure a WINS server or a master browser, typically on the same network as the security appliance, or reachable from that network. The WINS server or master browser provides the security appliance with an list of the resources on the network. You cannot use a DNS server instead.

Note **Note** File access is not supported in an Active Native Directory environment when used with Dynamic DNS. It is supported if used with a WINS server.

- **Enable file server entry**—Places the file server entry box on the portal page. File server access must be enabled.

  With this box selected, users can enter pathnames to Windows files directly. They can download, edit, delete, rename, and move files. They can also add files and folders. Again, shares must also be configured for user access on the applicable Windows servers. Users might have to be authenticated before accessing files, depending on network requirements.

- **Enable file server browsing**—Lets users browse the Windows network for domains/workgroups, servers and shares. File server access must be enabled.

  With this box checked, users can select domains and workgroups, and can browse servers and shares within those domains. Shares must also be configured for user access on the applicable Windows servers. Users may need to be authenticated before accessing servers, according to network requirements.

- **Enable port forwarding**—WebVPN Port Forwarding provides access for remote users in the group to client/server applications that communicate over known, fixed TCP/IP ports. Remote users can use client applications that are installed on their local PC and securely access a remote server that supports that application. Cisco has tested the following applications: Windows Terminal Services, Telnet, Secure FTP (FTP over SSH), Perforce, Outlook Express, and Lotus Notes. Other TCP-based applications may also work, but Cisco has not tested them.

Note Port Forwarding does not work with some SSL/TLS versions.

With this box checked users can access client/server applications by mapping TCP ports on the local and remote systems.

Note When users authenticate using digital certificates, the TCP Port Forwarding JAVA applet does not work. JAVA cannot access the web browser's keystore; therefore JAVA cannot use the certificates that the browser uses for user authentication, and the application cannot start. Do not use digital certificates to authenticate WebVPN users if you want them to be able to access applications.

- **Enable Outlook/Exchange proxy**—Enables the use of the Outlook/Exchange e-mail proxy.
- **Apply Web-type ACL**—Applies the WebVPN Access Control List defined for the users of this group.
- **Enable HTTP Proxy**—Enables the forwarding of an HTTP applet proxy to the client. The proxy is useful for technologies that interfere with proper mangling, such as Java, ActiveX, and Flash. It bypasses mangling while ensuring the continued use of the security appliance. The forwarded proxy modifies the browser's old proxy configuration automatically and redirects all HTTP and HTTPS

requests to the new proxy configuration. It supports virtually all client side technologies, including HTML, CSS, JavaScript, VBScript, ActiveX, and Java. The only browser it supports is Microsoft Internet Explorer.

- **Content Filtering**—Blocks or removes the parts of websites that use Java or Active X, scripts, display images, and deliver cookies. By default, these parameters are disabled, which means that no filtering occurs.
  - **Filter Java/ActiveX**—Removes <applet>, <embed> and <object> tags from HTML.
  - **Filter scripts**—Removes <script> tags from HTML.
  - **Filter images**—Removes <img> tags from HTML. Removing images dramatically speeds the delivery of web pages.
  - **Filter cookies from images**—Removes cookies that are delivered with images. This may preserve user privacy, because advertisers use cookies to track visitors.
- **Homepage**—Configures what, if any, home page to use.
  - **Specify URL**—Indicates whether the subsequent fields specify the protocol, either http or https, and the URL of the Web page to use as the home page.
  - **Protocol**—Specifies whether to use http or https as the connection protocol for the home page.
  - **://**—Specifies the URL of the Web page to use as the home page.
  - **Use none**—Specifies that no home page is configured.
- **Port Forwarding**—Configures port forwarding parameters.
  - **Port Forwarding List**—Specifies whether to inherit the port forwarding list from the default group policy, select one from the list, or create a new port forwarding list.
  - **New**—Displays a new panel on which you can add a new port forwarding list. See the description of the Add/Edit Port Forwarding List panel.
  - **Applet Name**—Specifies whether to inherit the applet name or to use the name specified in the box. Specify this name to identify port forwarding to end users. The name you configure displays in the end user interface as a hotlink. When users click this link, a Java applet opens a window that displays a table that lists and provides access to port forwarding applications that you configure for these users.The default applet name is Application Access.
- **Other**—Configures servers and URL lists and the Web-type ACL ID.
  - **Servers and URL Lists**—Specifies whether to inherit the list of Servers and URLs, to select and existing list, or to create a new list.
  - **New**—Displays a new panel on which you can add a new port forwarding list.
  - **Web-Type ACL ID**—Specifies the identifier of the Web-Type ACL to use.
- **SSL VPN Client tab**—lets you configure the security appliance to download SSL VPN clients (SVCs) to remote computer.

  SVC is a VPN tunneling technology that gives remote users the benefits of an IPSec VPN client without the need for network administrators to install and configure IPSec VPN clients on remote computers. The SVC uses the SSL encryption that is already present on the remote computer as well as the WebVPN login and authentication of the security appliance.

  To establish an SVC session, the remote user enters the IP address of a WebVPN interface of the security appliance in the browser, and the browser connects to that interface and displays the WebVPN login screen. If the user satisfies the login and authentication, and the security appliance identifies the user as *requiring* the SVC, the security appliance downloads the SVC to the remote

computer. If the security appliance identifies the user as having the *option* to use the SVC, the security appliance downloads the SVC to the remote computer while presenting a link on the user screen to skip the SVC installation.

After downloading, the SVC installs and configures itself, and then the SVC either remains or uninstalls itself (depending on the configuration) from the remote computer when the connection terminates.

The security appliance might have several unique SVC images residing in cache memory for different remote computer operating systems. When the user attempts to connect, the security appliance can consecutively download portions of these images to the remote computer until the image and operating system match, at which point it downloads the entire SVC. You can order the SVC images to minimize connection setup time, with the first image downloaded representing the most commonly-encountered remote computer operating system.

- **Inherit**—Indicates that the corresponding setting takes its value from the default group policy, rather than from the explicit specifications that follow.
- **Keep Installer on Client System**—Enables permanent SVC installation and disables the automatic uninstalling feature of the SVC. The SVC remains installed on the remote computer for subsequent SVC connections, reducing the SVC connection time for the remote user.
- **Keepalive Messages**—Adjusts the frequency of keepalive messages, in the range of 15 to 600 seconds. The default is keepalive messages are disabled.

  You can adjust the frequency of keepalive messages to ensure that an SVC connection through a proxy, firewall, or NAT device remains open, even if the device limits the time that the connection can be idle. Adjusting the frequency also ensures that the SVC does not disconnect and reconnect when the remote user is not actively running a socket-based application, such as Microsoft Outlook or Microsoft Internet Explorer.
- **Compression**—Enables compression on the SVC connection. By default, compression is enabled.

  SVC compression increases the communications performance between the security appliance and the SVC by reducing the size of the packets being transferred.
- **Rekey Negotiation Settings** group box—When the security appliance and the SVC perform a rekey, they renegotiate the crypto keys and initialization vectors, increasing the security of the connection.

  Renegotiation Interval specifies the number of minutes from the start of the session until the rekey takes place, from 1 to 10080 (1 week).

  Renegotiation Method specifies whether the SVC establishes a new tunnel during SVC rekey. If you check *none*, SVC rekey is disabled. If you check *ssl*, SSL renegotiation takes place during SVC rekey.
- **Dead Peer Detection**—Dead Peer Detection (DPD) ensures that the security appliance (gateway) or the SVC can quickly detect a condition where the peer is not responding, and the connection has failed.

  Gateway Side Detection enables DPD performed by the security appliance (gateway) and specifies the frequency, from 30 to 3600 seconds, with which the security appliance performs DPD. If you check *disable*, DPD performed by the security appliance is disabled.

  Client Side Detection enables DPD performed by the SVC (client), and specifies the frequency, from 30 to 3600 seconds, with which the SVC performs DPD. If you check *disable*, DPD performed by the SVC is disabled

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Auto Update

Configuration > Properties > Auto Update

The Auto Update pane lets you configure the security appliance to be managed remotely from servers that supports the Auto Update specification. Auto Update lets you apply configuration changes to the security appliance and receive software updates from remote locations.

Auto Update is useful in solving many of the challenges facing administrators for security appliance management:

- Overcomes dynamic addressing and NAT challenges.
- Gives ability to commit configuration changes in one atomic action.
- Provides a reliable method for updating software.
- Leverages well understood methods for high scalability.
- Open interface gives developers tremendous flexibility.
- Simplifies security solutions for Service Provider environments.
- High reliability, rich security/management features, broad support by many products.

Introduction to Auto Update

The Auto Update specification provides the infrastructure necessary for remote management applications to download security appliance configurations, software images, and to perform basic monitoring from a centralized location or multiple locations.

The Auto Update specification allows the Auto Update server to either push configuration information and send requests for information to the security appliance, or to pull configuration information by causing the security appliance to periodically poll the Auto Update server. The Auto Update server can also send a command to the security appliance to send an immediate polling request at any time. Communication between the Auto Update server and the security appliance requires a communications path and local CLI configuration on each security appliance.

The Auto Update feature on the security appliance can be used with Cisco security products, as well as products from third-party companies that want to manage the security appliance.

Important Notes

- If the security appliance configuration is updated from an Auto Update server, ASDM is not notified. You must choose **Refresh** or **File > Refresh ASDM with the Running Configuration on the Device** to get the latest configuration, and any changes to the configuration made in ASDM will be lost.

- If HTTPS is chosen as the protocol to communicate with the Auto Update server, the security appliance will use SSL. This requires the security appliance to have a DES or 3DES license.

Fields

The Auto Update pane consists of an Auto Update Servers table and two areas: the Timeout area, and the Polling area.

The Auto Update Servers table lets you view the parameters of previously-configured Auto Update servers. The security appliance polls the server listed at the top of the table first. You can change the position of the servers in the table with the Move Up and Move Down buttons. The Auto Update Servers table contains the following columns:

- Server—The name or IP address of the Auto Update server.
- User Name—The user name used to access the Auto Update server.
- Interface—The interface used when sending requests to the Auto Update server.
- Verify Certificate—Indicates whether the security appliance checks the certificate returned by the Auto Update server against the Certification Authority (CA) root certificates. This requires that the Auto Update server and the security appliance use the same CA.

Double-clicking any of the rows in the Auto Update Server table opens the Edit Auto Update Server dialog, in which you can modify the Auto Update server parameters. These changes are immediately reflected in the table, but you must click Apply to save them to the configuration.

The Timeout area lets you set the amount of time the security appliance waits for the Auto Update server to timeout. The Timeout area contains the following fields:

- Enable Timeout Period—Check to enable the security appliance to timeout if no response is received from the Auto Update server.
- Timeout Period (Minutes)—Enter the number of minutes the security appliance will wait to timeout if no response is received from the Auto Update server.

The Polling area lets you configure how often the security appliance will poll for information from the Auto Update server. The Polling area contains the following fields:

- Polling Period (minutes)—The number of minutes the security appliance will wait to poll the Auto Update server for new information.
- Poll on Specified Days—Allows you to specify a polling schedule.
- Set Polling Schedule—Displays the Set Polling Schedule dialog where you can configure the days and time-of-day to poll the Auto Update server.
- Retry Period (minutes)—The number of minutes the security appliance will wait to poll the Auto Update server for new information if the attempt to poll the server fails.
- Retry Count—The number of times the security appliance will attempt to retry to poll the Auto Update server for new information.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

Set Polling Schedule

Configuration > Properties > Auto Update > Set Polling Schedule

The Set Polling Schedule dialog lets you configure specific days, and the time-of-day for the security appliance to poll the Auto Update server.

Fields

The Set Polling Schedule dialog contains the following fields:

Days of the Week—Check the days of the week that you want the security appliance to poll the Auto Update server.

The Daily Update Window group lets you configure the time of day when you want the security appliance to poll the Auto Update server, and contains the following fields:

- Start Time—Enter the hour and minute to begin the Auto Update poll.
- Enable Randomize—Check to enable the security appliance to randomly choose a time to poll the Auto Update server.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

Add/Edit Auto Update Server

Configuration > Properties > Auto Update > Add/Edit Auto Update Server

The Edit Auto Update Server dialog contains the following fields:

- URL—The protocol the Auto Update server uses to communicate with the security appliance, either http or https, and the path to the Auto Update server.
- Interface—The interface to use when sending requests to the Auto Update server.
- Verify Certificate—Select to enable the security appliance to verify the certificate returned by the Auto Update server against the Certification Authority (CA) root certificates. This requires that the Auto Update server and the security appliance use the same CA.

The User area contains the following fields:

- User Name (Optional)—Enter the user name needed to access the Auto Update server.
- Password—Enter the user password for the Auto Update server.
- Confirm Password—Reenter the user password for the Auto Update server.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

Advanced Auto Update Settings

Configuration > Properties > Auto Update > Advanced Auto Update Settings

Fields

- Use Device ID to uniquely identify the ASA—Enables authentication using a Device ID. The Device ID is used to uniquely identify the security appliance to the Auto Update server.
- Device ID—Type of Device ID to use.
  - Hostname—The name of the host.
  - Serial Number—Device serial number.
  - IP Address on interface—The IP address of the selected interface, used to uniquely identify the security appliance to the Auto Update server.
  - MAC Address on interface—The MAC address of the selected interface, used to uniquely identify the security appliance to the Auto Update server.
  - User-defined value—A unique user ID.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

Client Update

Configuration > Properties > Client Update

The Client Update pane lets you configure the parameters of Auto Update clients associated with the security appliance when it is configured as an Auto Update server.

As an Auto Update server, you can specify the platform and asdm images for security appliances configured as Auto Update clients, including image revision numbers and locations, according to the device ID, device family, or device type of the client.

Introduction to Auto Update Server and Client Update

The Auto Update specification provides the infrastructure necessary for remote management applications to download security appliance configurations, software Images, and to perform basic monitoring from a centralized location.

As an Auto Update server, the specification allows the Auto Update server to either push configuration information and send requests for information to the security appliance, or to pull configuration information by causing the security appliance to periodically poll the Auto Update server. The Auto Update server can also send a command to the security appliance to send an immediate polling request at any time. Communication between the Auto Update server and the security appliance requires a communications path and local CLI configuration on each security appliance.

Fields

The Client Update pane consists of the following fields:

- Enable Client Update—Check to allow the security appliance to update the images used by other security appliances that are configured as Auto Update clients.
- Client Images table—lets you view previously-configured Client Update entries and includes the following columns:
  - Device—Displays a text string corresponding to a device-id of the client.
  - Device Family—Displays the family name of a client, either asa, pix, or a text string.
  - Device Type—Displays the type name of a client.
  - Image Type—Specifies the type of image, either ASDM image or Boot image.
  - Image URL—Specifies the URL for the software component.
  - Client Revision—Specifies the revision number(s) of the software component.

  Double-clicking any of the rows in the Client Images table opens the Edit Client Update Entry dialog, in which you can modify the client parameters. These changes are immediately reflected in the table, but you must click Apply to save them to the configuration.
- Live Client Update area—Lets you immediately update Auto Update clients that are currently connected to the security appliance through a tunnel.
  - Tunnel Group—Select “all” to update all Auto Update clients connected over all tunnel groups, or specify a tunnel group for clients that you want to update.
  - Update Now—Click to begin an immediate update.

Note Live Client Update is only available when the security appliance is configured in routed mode.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

Add/Edit Client Update

Configuration > Properties > Add/Edit Client Update

Fields

The Add/Edit Client Update dialog displays the following fields:

- Device Identification group:
  - Device ID—Enable if the client is configured to identify itself with a unique string, and specify the same string that the client uses. The maximum length is 63 characters.
  - Device Family—Enable if the client is configured to identify itself by device family, and specify the same same device family that the client uses. It can be asa, pix, or a text string with a maximum length of 7 characters.
  - Device Type—Enable if the client is configured to identify itself by device type, and specify the same device type that the client uses. It can be pix-515, pix-515e, pix-525, pix-535, asa5505, asa5510, asa5520, or asa5540. It can also be a text string with a maximum length of 15 characters.
  - Not Specified—Select for clients that do not match the above.
- Image Type—Specifies an image type, either ASDM or boot image. This URL must point to a file appropriate for this client. Maximum length of 255 characters.
- Client Revision—Specifies a text string corresponding to the revision number(s) of the software component. For example: 7.1(0)22.
- Image URL—Specifies the URL for the software component. This URL must point to a file appropriate for this client.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

DHCP and DNS Services

A DHCP server provides network configuration parameters, such as IP addresses, to DHCP clients. The security appliance can provide DHCP server or DHCP relay services to DHCP clients attached to security appliance interfaces. The DHCP server provides network configuration parameters directly to DHCP clients. DHCP relay passes DHCP requests received on one interface to an external DHCP server located behind a different interface.

The Domain Name System (DNS) is the system in the Internet that maps names of objects (usually host names) into IP numbers or other resource record values. The namespace of the Internet is divided into domains, and the responsibility for managing names within each domain is delegated, typically to systems within each domain. DNS client services allows you to specify DNS servers to which the security appliance sends DNS requests, request timeout period, and other parameters.

Dynamic DNS (DDNS) update integrates DNS with DHCP. The two protocols are complementary: DHCP centralizes and automates IP address allocation; DDNS update automatically records the association between assigned addresses and host names at pre-defined intervals. DDNS allows frequently changing address-host name associations to be updated frequently. Mobile hosts, for example, can then move freely on a network without user or administrator intervention.

For information about configuring these services, see the following topics:

DHCP Relay

Configuration > Properties > DHCP Services > DHCP Relay

The DHCP Relay pane lets you configure DHCP relay services on the security appliance. DHCP relay passes DHCP requests received on one interface to an external DHCP server located behind a different interface. To configure DHCP relay, you need to specify at least one DHCP relay server and then enable a DHCP relay agent on the interface receiving DHCP requests.

Restrictions

- You cannot enable a DHCP relay agent on an interface that has a DHCP relay server configured for it.
- The DHCP relay agent works only with external DHCP servers; it will not forward DHCP requests to a security appliance interface configured as a DHCP server.

Prerequisites

Before you can enable a DHCP relay agent on an interface, you must have at least one DHCP relay server in the configuration.

Fields

- DHCP Relay Agent—*Display only.* Contains the fields for configuring the DHCP relay agent.
  - Interface—Displays the interface ID. Double-clicking an interface opens the Edit DHCP Relay Agent Settings dialog box, where you can enable the DHCP relay agent and configure the relay agent parameters.
  - DHCP Relay Enabled—Indicates whether the DHCP relay agent is enabled on the interface. This column displays "Yes" if the DHCP relay agent is enabled or "No" if the DHCP relay agent is not enabled on the interface.
  - Set Route—Indicates whether the DHCP relay agent is configured to modify the default router address in the information returned from the DHCP server. This column display "Yes" if the DHCP relay agent is configured to change the default router address to the interface address or "No" if the DHCP relay agent does not modify the default router address.
  - Edit—Opens the Edit DHCP Relay Agent Settings dialog box, where you can enable the DHCP relay agent and configure the relay agent parameters.
- DHCP Relay Server—Contains the fields for configuring the DHCP relay servers.
  - Timeout—Specifies the amount of time, in seconds, allowed for DHCP address negotiation. Valid values range from 1 to 3600 seconds. The default value is 60 seconds.
  - Server—*Display only.* Displays the IP address of a configured, external DHCP server. Double-clicking a server address opens the DHCP Relay - Edit DHCP Server dialog box, where you can edit the DHCP relay server settings.
  - Interface—*Display only.* Display the interface the specified DHCP server is attached to.
  - Add—Opens the DHCP Relay - Add DHCP Server dialog box, where you can specify a new DHCP relay server. You can define up to 4 DHCP relay servers on the security appliance. This button is unavailable if you already have 4 DHCP relay servers defined.
  - Edit—Opens the DHCP Relay - Edit DHCP Server dialog box, where you can edit the DHCP relay server settings.
  - Delete—Removes the selected DHCP relay server. The server is removed from the security appliance configuration when you apply or save your changes.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Edit DHCP Relay Agent Settings

Configuration > Properties > DHCP Services > DHCP Relay > Edit DHCP Relay Agent Settings

You can enable the DHCP relay agent and configure the relay agent parameters for the selected interface in the Edit DHCP Relay Agent Settings dialog box.

Restrictions

- You cannot enable a DHCP relay agent on an interface that has a DHCP relay server configured for it.
- You cannot enable a DHCP relay agent on a security appliance that has DHCP server configured on an interface.

Prerequisites

Before you can enable a DHCP relay agent on an selected interface, you must have at least one DHCP relay server in the configuration.

Fields

- Enable DHCP Relay Agent—When checked, enables the DHCP relay agent on the selected interface. You must have a DHCP relay server defined before enabling the DHCP relay agent.
- Set Route—Specifies whether the DHCP relay agent is configured to modify the default router address in the information returned from the DHCP server. When this check box is checked, the DHCP relay agent substitutes the address of the selected interface for the default router address in the information returned from the DHCP server.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

DHCP Relay - Add/Edit DHCP Server

Configuration > Properties > DHCP Services > DHCP Relay > DHCP Relay - Add/Edit DHCP Server

Define new DHCP relay servers in the DHCP Relay - Add DHCP Server dialog box or edit exiting server information in the DHCP Relay - Edit DHCP Server dialog box. You can define up to 4 DHCP relay servers.

Restrictions

You cannot define a DHCP relay server on an interface with a DHCP server enabled on it.

Fields

- DHCP Server—Specifies the IP address of the external DHCP server to which DHCP requests are forwarded.
- Interface—Specifies the interface through which DHCP requests are forwarded to the external DHCP server.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

DHCP Server

Configuration > Properties > DHCP Services > DHCP Server

The DHCP Server pane lets you configure the security appliance interfaces as DHCP servers. You can configure one DHCP server per interface on the security appliance.

Note You cannot configure a DHCP server on an interface that has DHCP relay configured on it. For more information about DHCP relay, see DHCP Relay.

Fields

- Interface—*Display only.* Displays the interface ID. Double-clicking an interface ID opens the Edit DHCP Server dialog box, where you can enable DHCP on and assign a DHCP address pool to the selected interface.
- DHCP Enabled—*Display only.* Indicates whether DHCP is enabled on the interface. This column displays “Yes” if DHCP is enabled or “No” if DHCP is not enabled on the interface.
- Address Pool—*Display only.* Displays the range of IP addresses assigned to the DHCP address pool.
- DNS Servers—*Display only.* Displays the DNS servers configured for the interface.
- WINS Servers—*Display only.* Displays the WINS servers configured for the interface.
- Domain Name—*Display only.* Displays the domain name of the interface.
- Ping Timeout—*Display only.* Displays time in milliseconds that the security appliance will wait for an ICMP ping response on the interface.
- Lease Length—*Display only.* Displays the duration of time that the DHPC server configured on the interface allows DHCP clients to use the an assigned IP address.
- Auto Interface—*Display only.* Displays the interface on a DHCP client providing DNS, WINS, and domain name information for automatic configuration.
- Options—*Display only.* Displays advanced DHCP options configured for the interface.
- Dynamic DNS Settings—*Display only.* Displays
- Edit—Opens the Edit DHCP Server dialog box for the selected interface. You can enable DHCP and specify the DHCP address pool in the Edit DHCP Server dialog box.
- Other DHCP Options—Contains optional DHCP parameters.
  - Enable Autoconfiguration on interface—Check to enable DHCP auto configuration and select the interface from the menu.

DHCP auto configuration causes the DHCP server to provide DHCP clients with DNS server, domain name, and WINS server information obtained from a DHCP client running on the specified interface. If any of the information obtained through auto configuration is also specified manually in the Other DHCP Options area, the manually specified information takes precedence over the discovered information.

- DNS Server 1—(Optional) Specifies the IP address of the primary DNS server for a DHCP client.
- DNS Server 2—(Optional) Specifies the IP address of the alternate DNS server for a DHCP client.
- Domain Name—(Optional) Specifies the DNS domain name for DHCP clients. Enter a valid DNS domain name, for example example.com.
- Lease Length—(Optional) Specifies the amount of time, in seconds, that the client can use its allocated IP address before the lease expires. Valid values range from 300 to 1048575 seconds. The default value is 3600 seconds (1 hour).
- Primary WINS Server—(Optional) Specifies the IP address of the primary WINS server for a DHCP client.
- Secondary WINS Server—(Optional) Specifies the IP address of the alternate WINS server for a DHCP client.
- Ping Timeout—(Optional) To avoid address conflicts, the security appliance sends two ICMP ping packets to an address before assigning that address to a DHCP client. The Ping Timeout field specifies the amount of time, in milliseconds, that the security appliance waits to time out a DHCP ping attempt. Valid values range from 10 to 10000 milliseconds. The default value is 50 milliseconds.
- Advanced—Opens the Advanced DHCP Options dialog box, where you can specify DHCP options and their parameters.

- Dynamic DNS Settings for DHCP Server—In this area, you can configure the DDNS update settings for the DHCP server.
  - Update DNS Clients—Check to specify that, besides the default action of updating the client PTR resource records, the DHCP server should also perform the following update actions (if selected):
  - Update Both Records—Check to specify that the DHCP server should update both the A and PTR RRs.
  - Override Client Settings—Check to specify that the DHCP server actions should override any update actions requested by the DHCP client.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Edit DHCP Server

Configuration > Properties > DHCP Services > DHCP Server > Edit DHCP Server

You can enable DHCP and specify the DHCP address pool for the selected interface in the Edit DHCP Server dialog box.

Fields

- Enable DHCP Server—Check this check box to enable the DHCP server on the selected interface. Uncheck this check box to disable DHCP on the selected interface. Disabling the DHCP server on the selected interface does not clear the specified DHCP address pool.
- DHCP Address Pool—Enter the IP address pool used by the DHCP server. Enter the range of IP addresses from lowest to highest. The range of IP addresses must be on the same subnet as the selected interface and cannot contain the IP address of the interface itself.
- Optional Parameters—You can optionally configure the following parameters for the DHCP server:
  - DNS Server 1—Enter the IP address of the primary DNS server for a DHCP client.
  - DNS Server 2— Enter the IP address of the alternate DNS server for a DHCP client.
  - Domain Name—Enter the DNS domain name for DHCP clients. Enter a valid DNS domain name, for example example.com.
  - Lease Length—Enter the amount of time, in seconds, that the client can use its allocated IP address before the lease expires. Valid values range from 300 to 1048575 seconds. The default value is 3600 seconds (1 hour).
  - Primary WINS Server—Enter the IP address of the primary WINS server for a DHCP client.
  - Secondary WINS Server—Enter the IP address of the alternate WINS server for a DHCP client.
  - Ping Timeout—Enter the amount of time, in milliseconds, that the security appliance waits to time out a DHCP ping attempt. Valid values range from 10 to 10000 milliseconds. The default value is 50 milliseconds.
  - Enable Autoconfiguration on interface—Check to enable DHCP auto configuration and select the interface from the menu.
  - Advanced—Opens the Advanced DHCP Options dialog box, where you can specify DHCP options and their parameters.
- Dynamic DNS Settings for DHCP Server—In this area, you can configure the DDNS update settings for the DHCP server.
  - Update DNS Clients—Check to specify that, besides the default action of updating the client PTR resource records, the DHCP server should also perform the following update actions (if selected):
  - Update Both Records—Check to specify that the DHCP server should update both the A and PTR RRs.
  - Override Client Settings—Check to specify that DHCP server actions should override any update actions requested by the DHCP client.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Advanced DHCP Options

Configuration > Properties > DHCP Services > DHCP Server > Advanced DHCP Options

The Advanced DHCP Options dialog box lets you configure DHCP option parameters. You use DHCP options to provide additional information to DHCP clients. For example, DHCP option 150 and DHCP option 66 provide TFTP server information to Cisco IP Phones and Cisco IOS routers.

You can use that advanced DHCP options to provide DNS, WINS, and domain name parameters to DHCP clients. You can also use the DHCP auto configuration setting to obtain these values or manually specify them on the DHCP Server pane. When you use more than one method to specify this information, the information is passed to DHCP clients with the following preference:

1. Manually configured settings.
2. Advanced DHCP Options settings.
3. DHCP auto configuration.

For example, you can manually define the domain name that you want the DHCP clients to receive, and then enable DHCP auto configuration. Although DHCP auto configuration will discover the domain along with the DNS and WINS servers, the manually-defined domain name is passed to DHCP clients with the discovered DNS and WINS server names. The domain name discovered by the DHCP auto configuration process is discarded in favor of the manually-defined domain name.

Fields

- Option to be Added—Contains the fields used to configure a DHCP option.
  - Choose the option code—Lists the available option codes. All DHCP options (options 1 through 255) are supported except 1, 12, 50–54, 58–59, 61, 67, and 82. Choose the option that you want to configure.

    Some options are standard. For standard options, the option name is shown in parentheses after the option number and the option parameters are limited to those supported by the option. For all other options, only the option number is shown and you must choose the appropriate parameters to supply with the option.

    For standard DHCP options, only the supported option value type is available. For example, if you choose DHCP Option 2 (Time Offset), you can only supply a hexadecimal value for the option. For all other DHCP options, all of the option value types are available and you must choose the appropriate options value type.
- Option Data—These options specify the type of information the option returns to the DHCP client. For standard DHCP options, only the supported option value type is available. For all other DHCP options, all of the option value types are available.
- IP Address—Choosing this value specifies that an IP address is returned to the DHCP client. You can specify up to two IP addresses.

Note The name of the associated IP Address fields can change based on the DHCP option you chose. For example, if you choose DHCP Option 3 (Router), the fields change name to Router 1 and Router 2.

  - IP Address 1—An IP address in dotted-decimal notation.
  - IP Address 2—(Optional) An IP address in dotted-decimal notation.
- ASCII—Choose this option specifies that an ASCII value is returned to the DHCP client.

Note The name of the associated Data field can change based on the DHCP option you chose. For example, if you choose DHCP Option 14 (Merit Dump File), the associated Data field changes name to File Name.

  - Data—An ASCII character string. The string cannot include white space.
- Hex—Selecting this option specifies that a hexadecimal value is returned to the DHCP client.

Note The name of the associated Data field can change based on the DHCP option you chose. For example, if you choose DHCP Option 2 (Time Offset), the associated Data field becomes the Offset field.

  - Data—A hexadecimal string with an even number of digits and no spaces. You do not need to use a 0x prefix.
- Add—Adds the configured option to the DHCP option table.
- Delete—Removes the selected option from the DHCP option table.
- DHCP option table—Lists the DHCP options that have been configured.
  - Option Code—Shows the DHCP option code. For standard DHCP options, the option name appears in parentheses next to the option code.
  - Option Data—Shows the parameters that have been configured for the selected option.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

DNS Client

Configuration > Properties > DNS > DNS Client

The DNS Client pane shows the DNS server groups and DNS lookup information for the security appliance, so it can resolve server names to IP addresses in your WebVPN configuration or certificate configuration. Other features that define server names (such as AAA) do not support DNS resolution. In those cases, you must enter the IP address or manually resolve the name to an IP address by adding the server name in the Network Object Groups pane.

Fields

- DNS Server Groups—Displays and manages the DNS server list. There can be up to six addresses to which DNS requests can be forwarded. The security appliance tries each DNS server in order until it receives a response. You must enable DNS on at least one interface in the DNS Lookup area before you can add a DNS server. The contents of the table in this area are as follows:
  - Name—*Display only.* Shows the name of each configured DNS server group.
  - Servers—*Display only.* Shows the IP addresses of the configured servers.
  - Timeout—*Display only.* Shows the number of seconds to wait before trying the next DNS server in the list, between 1 and 30 seconds. The default is 2 seconds. Each time the security appliance retries the list of servers, this timeout doubles.
  - Retries—*Display only.* Shows the number of seconds to wait before trying the next DNS server in the list.
  - Domain Name—*Display only.* Shows the number of times the security appliance retries the request.
- DNS Lookup—Enables or disables DNS lookup on an interface.
  - Interface—*Display only.* Lists all interface names.
  - DNS Enabled—*Display only.* Shows whether an interface supports DNS lookup, Yes or No.
  - Disable—Disables DNS lookup for the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit DNS Server Group

Configuration > Properties > DNS Client > Add/Edit DNS Server Group

The Add or Edit DNS Server Group pane lets you specify or modify one or more DNS servers for the security appliance so it can resolve server names to IP addresses in your WebVPN configuration or certificate configuration (See Add/Edit Trustpoint Configuration > Enrollment Settings Tab and

Add/Edit Trustpoint Configuration > CRL Retrieval Policy Tab). Other features that define server names (such as AAA) do not support DNS resolution. For those, you must enter the IP address or manually resolve the name to an IP address by adding the server name in the Network Object Groups pane.

Fields

- Name—Specifies the server name. For the Edit function, this field is *Display only.*
- DNS Servers—Manages the DNS server list. You can specify up to six addresses to which DNS requests can be forwarded. The security appliance tries each DNS server in order until it receives a response. You must enable DNS on at least one interface in the DNS Lookup area before you can add a DNS server.
  - Server to be Added—Specifies the DNS server IP address.
  - Add—Adds a DNS server to the bottom of the list.
  - Delete—Deletes the selected DNS server from the list.
  - Servers—*Display only.* Shows the DNS server list.
  - Move Up—Moves the selected DNS server up the list.
  - Move down—Moves the selected DNS server down the list.
- Timeout—Specifies the number of seconds to wait before trying the next DNS server in the list, between 1 and 30 seconds. The default is 2 seconds. Each time the security appliance retries the list of servers, this timeout doubles.
- Retries—Sets the number of times the security appliance retries the request. The range is 1 through 10 retries.
- Domain Name—(Optional) Specifies the DNS domain name for the server. Enter a valid DNS domain name; for example example.com.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Dynamic DNS

Configuration > Properties > DNS > Dynamic DNS

Dynamic DNS provides address and domain name mappings so hosts can find each other even though their DHCP-assigned IP addresses change frequently. The DDNS name and address mappings are held on the DHCP server in two resource records: the A RR contains the name to IP address mapping while the PTR RR maps addresses to names. Of the two methods for performing DDNS updates—the IETF standard defined by RFC 2136 and a generic HTTP method—the security appliance supports the IETF method in this release.

The Dynamic DNS pane shows the configured DDNS update methods and the interfaces configured for DDNS. By automatically records the association between assigned addresses and host names at pre-defined intervals, DDNS allows frequently changing address-host name associations to be updated frequently. Mobile hosts, for example, can then move freely on a network without user or administrator intervention.

Fields

- Update Methods—Lists the DDNS update methods that are configured on the security appliance. This table includes:
  - Method Name—*Display only.* Shows the user-defined name for the DDNS update method.
  - Interval—*Display only.* Shows the time between DNS update attempts configured for the update method.
  - Update DNS Server Records—*Display only.* Shows whether the method updates both the A resource record (name to IP address) and the PTR resource record (IP address to name), or neither record.
  - Add/Edit—Displays the Add/Edit Dynamic DNS Update Methods dialog box.
  - Delete—Removes the currently selected update method from the table.
- Dynamic DNS Interface Settings—Lists the DDNS settings for each interface configured for DDNS.
  - Interface—*Display only.* Shows the names of the security appliance interfaces configured for DDNS.
  - Method Name—*Display only.* Shows the update methods assigned to each interface.
  - Hostname—*Display only.* Shows the hostname of the DDNS client.
  - Update DHCP Server Records—*Display only.* Shows whether the interface updates both the A and PTR resource records or neither.
  - Add/Edit—Displays the Add/Edit Dynamic DNS Interface Settings dialog box.
  - Delete—Removes the DDNS update settings for the selected interface.
- DHCP Clients Update DNS Records—This is the global setting specifying which records the DHCP client requests to be updated by the DHCP server. Click one of the following radio buttons:
  - Default (PTR Records) to specify that the client request PTR record updating by the server

    –or–
  - Both (PTR Records and A Records) to specify that the client request both the A and PTR DNS resource records by the server

    –or–
  - None to specify that the client request no updates by the server

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

Add/Edit Dynamic DNS Update Methods

Configuration > Properties > DNS > Dynamic DNS > Add/Edit Dynamic DNS Update Method

The Add/Edit Dynamic DNS Update Methods dialog box lets you add a new method or edit a previously added method. You can specify the method name (if adding a method), specify the interval between DDNS update attempts, and specify whether the DDNS client attempts to update both or neither of the two DNS records, the A record and the PTR record.

Fields

- Name—If you are adding a method, enter then name of the new method in this field. If you are editing an existing method, this field is *display-only* and shows the name of the method selected for editing.
- Update Interval—Specifies the time to elapse between update attempts. The interval ranges from 0 to nearly one year.
  - Days—Choose the number of days between update attempts from 0 to 364.
  - Hours—Choose the number of hours (in whole numbers) between update attempts from 0 to 23.
  - Minutes—Choose the number of minutes (in whole numbers) between update attempts from 0 to 59.
  - Seconds—Choose the number of minutes (in whole numbers) between update attempts from 0 to 59.
  - Update Records—Click Both (A and PTR Records) for the client to attempt updates to both the A and PTR DNS resource records, or click A Records Only to update just the A records. This is the individual method setting for DNS server records updated by the client.

These units are additive. That is, if you enter 0 days, 0 hours, 5 minutes and 15 seconds, the update method will attempt an update every 5 minutes and 15 seconds for as long as the method is active.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

Add/Edit Dynamic DNS Interface Settings

Configuration > Properties > DNS > Dynamic DNS > Add/Edit Dynamic DNS Interface Settings

The Add/Edit Dynamic DNS Interface Settings allows you to configure DDNS on a security appliance interface. You can assign an update method, specify the hostname, and configure DHCP server updating of both the A and PTR records by the client or neither.

Fields

- Interface—Choose an interface on which to configure DDNS from the menu.
- Update Method—Choose an available DDNS update method from the menu.

- Hostname—Enter the hostname of the DDNS client.
- DHCP Client—This area allows you to specify that the DHCP client updates both the A and PTR DNS records or neither. This interface setting overrides the global setting at Configuration > Properties > DNS > Dynamic DNS
- DCHP Client Updates DNS Records—Click one of the following radio buttons:
  - Default (PTR Records only) to specify that the client request only PTR record updating by the server

    –or–
  - Both (PTR Records and A Records) to specify that the client request both the A and PTR DNS resource records by the server

    –or–
  - None to specify that the client request no updates by the server

Note DHCP must be enabled on the selected interface for this action to be effective.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

Configuring AAA Servers

This section contains the following topics:

Understanding AAA

This section contains the following topics:

AAA Overview

AAA enables the security appliance to determine who the user is (authentication), what the user can do (authorization), and what the user did (accounting). You can use authentication alone or with authorization and accounting. Authorization always requires a user to be authenticated first. You can use accounting alone, or with authentication and authorization.

AAA provides an extra level of protection and control for user access than using access lists alone. For example, you can create an access list allowing all outside users to access Telnet on a server on the DMZ network. If you want only some users to access the server and you might not always know IP addresses of these users, you can enable AAA to allow only authenticated and/or authorized users to make it through the security appliance. (The Telnet server enforces authentication, too; the security appliance prevents unauthorized users from attempting to access the server.)

- About Authentication—Authentication grants access based on user identity. Authentication establishes user identity by requiring valid user credentials, which are typically a username and password.
- About Authorization—Authorization controls access per user after users authenticate. Authorization controls the services and commands available to each authenticated user. Were you not to enable authorization, authentication alone would provide the same access to services for all authenticated users.

If you need the control that authorization provides, you can configure a broad authentication rule, and then have a detailed authorization configuration. For example, you authenticate inside users who attempt to access any server on the outside network and then limit the outside servers that a particular user can access using authorization.

The security appliance caches the first 16 authorization requests per user, so if the user accesses the same services during the current authentication session, the security appliance does not resend the request to the authorization server.

- About Accounting—Accounting tracks traffic that passes through the security appliance, enabling you to have a record of user activity. If you enable authentication for that traffic, you can account for traffic per user. If you do not authenticate the traffic, you can account for traffic per IP address. Accounting information includes when sessions start and stop, username, the number of bytes that pass through the security appliance for the session, the service used, and the duration of each session.

Preparing for AAA

AAA services depend upon the use of the LOCAL database or at least one AAA server. You can also use the LOCAL database as a fallback for most services provided by a AAA server. Before you implement AAA, you should configure the LOCAL database and configure AAA server groups and servers.

How you configure the LOCAL database and AAA servers depends upon the AAA services you want the security appliance to support. Regardless of whether you use AAA servers, you should configure the LOCAL database with user accounts that support administrative access, to prevent accidental lockouts and, if so desired, to provide a fallback method when AAA servers are unreachable. For more information, see LOCAL Database.

Table 10-1 provides a summary of AAA service support by each AAA server type and by the LOCAL database. You manage the LOCAL database by configuring user profiles in the Configuration > Properties > Device Administration > User Accounts pane. You establish AAA server groups and add individual AAA servers to the server groups in the Configuration > Properties > AAA Setup > AAA Server Groups pane.

Table 10-1 Summary of AAA Support

	Database Type							
AAA Service	**Local**	**RADIUS**	**TACACS+**	**SDI**	**NT**	**Kerberos**	**LDAP**	**HTTP Form**
Authentication of...								
VPN users	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes[1]
Firewall sessions	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
Administrators	Yes	Yes	Yes	Yes[2]	Yes	Yes	Yes	No
Authorization of...								
VPN users	Yes	Yes	No	No	No	No	Yes	No
Firewall sessions	No	Yes[3]	Yes	No	No	No	No	No
Administrators	Yes[4]	No	Yes	No	No	No	No	No
Accounting of...								
VPN connections	No	Yes	Yes	No	No	No	No	No

Table 10-1 Summary of AAA Support (continued)

AAA Service	Database Type							
	Local	RADIUS	TACACS+	SDI	NT	Kerberos	LDAP	HTTP Form
Firewall sessions	No	Yes	Yes	No	No	No	No	No
Administrators	No	Yes[5]	Yes	No	No	No	No	No

1. HTTP Form protocol supports single sign-on authentication for WebVPN users only.
2. SDI is not supported for HTTP administrative access.
3. For firewall sessions, RADIUS authorization is supported with user-specific access lists only, which are received or specified in a RADIUS authentication response.
4. Local command authorization is supported by privilege level only.
5. Command accounting is available for TACACS+ only.

LOCAL Database

The security appliance maintains a local database that you can populate with user profiles.

- User Profiles—User profiles contain, at a minimum, a username. Typically, you assign a password to each username, although passwords are optional. User profiles can also specify VPN access policy per user. You can manage user profiles with the Configuration > Properties > Device Administration > User Accounts pane.
- Fallback Support—The local database can act as a fallback method for console and enable password authentication, for command authorization, and for VPN authentication and authorization. This behavior is designed to help you prevent accidental lockout from the security appliance. For users who need fallback support, we recommend that their usernames and passwords in the local database match their usernames and passwords in the AAA servers. This provides transparent fallback support. Because the user cannot determine whether a AAA server or the local database is providing the service, using usernames and passwords on AAA servers that are different than the usernames and passwords in the local database means that the user cannot be certain which username and password should be given.

AAA Implementation in ASDM

You can use AAA for the following:

- AAA for Device Administration
- AAA for Network Access
- AAA for VPN Access

AAA for Device Administration

You can authenticate all administrative connections to the security appliance, including:

- Telnet
- SSH
- Serial console

- ASDM
- VPN management access

You can also authenticate administrators who attempt to enter enable mode. You can authorize administrative commands. You can have accounting data for administrative sessions and for commands issued during a session sent to an accounting server.

You can configure AAA for device administration with the Configuration > Properties > Device Access > AAA Access pane.

AAA for Network Access

You can configure rules for authenticating, authorizing, and accounting for traffic passing through the firewall by using the Configuration > Security Policy > AAA Rules tab. The rules you create are similar to access rules, except that they specify whether to authenticate, authorize, or perform accounting for the traffic defined; and which AAA server group the security applianceis to use to process the AAA service request.

AAA for VPN Access

AAA services for VPN access include the following:

- User account settings for assigning users to VPN groups, configured in the Configuration > Properties > Device Administration > User Accounts pane.
- VPN group policies that can be referenced by many user accounts or tunnel groups, configured in the Configuration > VPN > General > Group Policy pane.
- Tunnel group policies, configured in the Configuration > VPN > General > Tunnel Group pane.

AAA Setup

The AAA Setup panes let you configure AAA server groups, AAA servers, and the authentication prompt. This section includes the following topics:

AAA Server Groups

Configuration > Properties > AAA Setup > AAA Server Groups

The AAA Server Groups pane lets you:

- Configure AAA server groups and the protocols the security appliance uses to communicate with the servers listed in each group.
- Configure and add individual servers to AAA server groups.

You can have up to 15 groups in single-mode or 4 groups in multi-mode. Each group can have up to 16 servers in single mode or 4 servers in multi-mode. When a user logs in, the servers are accessed one at a time, starting with the first server you specify, until a server responds.

If AAA accounting is in effect, the accounting information goes only to the active server, unless you have configured simultaneous accounting.

For an overview of AAA services, see Understanding AAA.

Fields

The fields in the AAA Server Groups pane are grouped into two main areas: the AAA Server Groups area and the Servers In The Selected Group area. The AAA Server Groups area lets you configure AAA server groups and the protocols the security appliance uses to communicate with the servers listed in each group.

Note Double-clicking any of the rows in the AAA Server Groups table opens the Edit AAA Server Group dialog box, in which you can modify the AAA Server Group parameters. These changes are immediately reflected in the table, but you must click **Apply** to save them to the configuration.

Clicking a column head sorts the table rows in alphanumeric order according to the contents of that column.

- Server Group—*Display only.* Shows the symbolic name of the selected server group.
- Protocol—*Display only.* Lists the AAA protocol that servers in the group support.
- Accounting Mode—*Display only.* Shows either simultaneous or single mode accounting. In single mode, the security appliance sends accounting data to only one server. In simultaneous mode, the security appliance sends accounting data to all servers in the group.
- Reactivation Mode—*Display only.* Shows the method by which failed servers are reactivated: Depletion or Timed reactivation mode. In Depletion mode, failed servers are reactivated only after all of the servers in the group are inactive. In Timed mode, failed servers are reactivated after 30 seconds of down time.
- Dead Time—*Display only.* Shows the number of minutes that will elapse between the disabling of the last server in the group and the subsequent reenabling of all servers. This parameter applies only in depletion mode.
- Max Failed Attempts—*Display only.* Shows the number of failed connection attempts allowed before declaring a nonresponsive server inactive.
- Add—Displays the Add AAA Server Group dialog box.
- Edit—Displays the Edit AAA Server Group dialog box, or, if you have selected LOCAL as the server group, displays the Edit AAA Local Server Group dialog box.
- Delete—Removes the currently selected server group entry from the server group table. There is no confirmation or undo.

The Servers In Selected Group area, the second area of the AAA Server Groups pane, lets you add and configure AAA servers for existing AAA server groups. The servers can be RADIUS, TACACS+, NT, SDI, Kerberos, LDAP, or HTTP-form servers.

- Server Name or IP Address—*Display only.* Shows the name or IP address of the AAA server.
- Interface—*Display only.* Shows the network interface where the authentication server resides.

- Timeout—*Display only.* Shows the timeout interval, in seconds. This is the time after which the security appliance gives up on the request to the primary AAA server. If there is a standby AAA server, the security appliance sends the request to the backup server.
- Add/Edit—Displays the Add/Edit AAA Server dialog box.
- Delete—Removes the selected AAA server from the list.
- Move up—Moves the selected AAA server up in the AAA sequence.
- Move down—Moves the selected AAA server back in the AAA sequence.
- Test—Displays the Test AAA Server dialog box.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	• [1]	•	•	—

1. HTTP Form and WebVPN are supported only in single routed mode.

Add/Edit AAA Server Group

Configuration > Properties > AAA Setup > AAA Server Groups > Add/Edit AAA Server Group

The Add/Edit AAA Server Group dialog box lets you add or modify AAA server groups. The results appear in the AAA Server table.

Fields

- Server Group—*Display only.* Shows the name of the selected server group.
- Protocol drop-down list—Specifies the protocols supported by servers in the group. They include RADIUS, TACACS+, NT Domain, SDI, Kerberos, LDAP, and HTTP Form for single sign-on (WebVPN users only).

Note The following fields are not available after selecting the HTTP Form protocol.

- Accounting Mode—Specifies the accounting mode used with the server group.
  - Simultaneous—Configures the security appliance to send accounting data to all servers in the group.
  - Single—Configures the security appliance to send accounting data to only one server of the group.
- Reactivation Mode—Specifies the method by which failed servers are reactivated.
  - Depletion—Configures the security appliance to reactivate failed servers only after all of the servers in the group are inactive.
  - Timed—Configures the security appliance to reactive failed servers after 30 seconds of down time.

- Dead Time—Specifies the number of minutes that will elapse between the disabling of the last server in the group and the subsequent reenabling of all servers. This field is not available for timed mode.
- Max Failed Attempts—Specifies the number of failed connection attempts (1 through 5) allowed before declaring a nonresponsive server inactive.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	• [1]	•	•	—

1. HTTP Form and WebVPN are supported only in single routed mode.

Edit AAA Local Server Group

Configuration > Properties > AAA Setup > AAA Server Groups > Edit AAA Local Server Group

The Edit AAA Local Server Group dialog box lets you specify whether to enable local user lockout and the maximum number of failed login attempts to allow before locking out the user. If a user is locked out, and administrator must clear the lockout condition before the user can successfully log in.

Fields

- Enable Local User Lockout —Enables locking out and denying access to a user who has exceeded the configured maximum number of failed authentication attempts.
- Maximum Attempts—Specifies the maximum number of failed login attempts allowed before locking out and denying access to a user. This limit applies only when the LOCAL database is used for authentication.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	• [1]	•	•	—

1. HTTP Form and WebVPN are supported only in single routed mode.

Add/Edit AAA Server

Configuration > Properties > AAA Setup > AAA Servers > Add/Edit AAA Server

The Add/Edit AAA Server dialog box lets you modify the parameters of an existing AAA server or add a new AAA server to an existing group selected in the AAA server groups table.

Fields

Note The first four fields are the same for all types of servers. The area contents area is specific to each server type.

- Server Group—*Display only.* Shows the name of the server group.
- Interface Name—Specifies the network interface where the server resides.
- Server Name or IP Address—Specifies the name or IP address of the AAA server.
- Timeout—Specifies the timeout interval, in seconds. This is the time after which the security appliance gives up on the request to the primary AAA server. If there is a standby AAA server, the security appliance sends the request to the backup server.
- RADIUS Parameters area—Specifies the parameters needed for using a RADIUS server. This area appears only when the selected server group uses RADIUS.
  - Retry Interval—Specifies the number of seconds to wait after sending a query to the server and receiving no response, before reattempting the connection. The minimum time is 1 second. The default time is 10 seconds. The maximum time is 10 seconds.
  - Server Authentication Port—Specifies the server port to use for user authentication. The default port is 1645.

Note The latest RFC states that RADIUS should be on UDP port number 1812, so you might need to change this default value to 1812.

  - Server Accounting Port—Specifies the server port to use for user accounting. The default port is 1646.
  - Server Secret Key—Specifies the server secret key (also called the shared secret) to use for encryption; for example: C8z077f. The secret is case-sensitive. The field displays only asterisks.The security appliance uses the server secret to authenticate to the RADIUS server. The server secret you configure here should match the one configured on the RADIUS server. If you do not know the server secret for the RADIUS server, ask the administrator of the RADIUS server. The maximum field length is 64 characters.
  - Confirm Server Secret Key—Requires that you reenter the server secret, to confirm its accuracy. The secret is case-sensitive. The field displays only asterisks.
  - Common Password—Specifies the common password for the group. The password is case-sensitive. The field displays only asterisks. If you are defining a RADIUS server to be used for authentication rather than authorization, do not provide a common password.

    A RADIUS authorization server requires a password and username for each connecting user. You enter the password here. The RADIUS authorization server administrator must configure the RADIUS server to associate this password with each user authorizing to the server via this security appliance. Be sure to provide this information to your RADIUS server administrator. Enter a common password for all users who are accessing this RADIUS authorization server through this security appliance.

    If you leave this field blank, each user password will be his or her own username. For example, a user with the username “jsmith” would enter “jsmith”. As a security precaution never use a RADIUS authorization server for authentication. Use of a common password or usernames as passwords is much less secure than strong passwords per user.

Note The password field is required by the RADIUS protocol and the RADIUS server requires it; however, users do not need to know it.

  - Confirm Common Password—Requires that you reenter the common password, to confirm its accuracy. The password is case-sensitive. The field displays only asterisks.
  - ACL Netmask Convert—Specifies how the security appliance handles netmasks received in downloadable access lists. The security appliance expects downloadable access lists to contain standard netmask expressions whereas Cisco Secure VPN 3000 series concentrators expect downloadable access lists to contain wildcard netmask expressions, which are the reverse of a standard netmask expression. A wildcard mask has ones in bit positions to ignore, zeros in bit positions to match. The ACL Netmask Convert list helps minimize the effects of these differences upon how you configure downloadable access lists on your RADIUS servers.

    If you choose Detect Automatically, the security appliance attempts to determine the type of netmask expression used. If it detects a wildcard netmask expression, it converts it to a standard netmask expression; however, because some wildcard expressions are difficult to detect unambiguously, this setting may occasionally misinterpret a wildcard netmask expression as a standard netmask expression.

    If you choose Standard, the security appliance assumes downloadable access lists received from the RADIUS server contain only standard netmask expressions. No translation from wildcard netmask expressions is performed.

    If you choose Wildcard, the security appliance assumes downloadable access lists received from the RADIUS server contain only wildcard netmask expressions and it converts them all to standard netmask expressions when the access lists are downloaded.

- TACACS+ Parameters—Specifies the parameters needed for using a TACACS+ server. This area appears only when the selected server group uses TACACS+.
  - Server Port—Specifies the server port to use.
  - Server Secret Key—Specifies the server secret key to use for encryption. The secret is case-sensitive. The field displays only asterisks.
  - Confirm Server Secret Key—Requires that you reenter the server secret, to confirm its accuracy. The secret is case-sensitive. The field displays only asterisks.
- SDI Parameters—Specifies the parameters needed for using an SDI server. This area appears only when the selected server group uses SDI.
  - Server Port—Specifies the server port to use.
  - Retry Interval—Specifies the number of seconds to wait before reattempting the connection.
- Kerberos Parameters—Specifies the parameters needed for using a Kerberos server. This area appears only when the selected server group uses Kerberos.
  - Server Port—Specifies the server port that the Kerberos server listens to.
  - Retry Interval—Specifies the number of seconds to wait before reattempting the connection. Enter the number of times to retry sending a query to the server after the timeout period. If there is still no response after this number of retries, the security appliance declares this server inoperative and uses the next Kerberos/Active Directory server in the list. The minimum number of retries is 0. The default number of retries is 2. The maximum number of retries is 10.

  - Kerberos Realm—Specifies the name of the Kerberos realm to use, for example: USDOMAIN.ACME.COM. The maximum length is 64 characters. The following types of servers require that you enter the realm name in all uppercase letters: Windows 2000, Windows XP, and Windows.NET. You must enter this name, and it must be the correct realm name for the server whose IP address you entered in the Server IP Address field.

- LDAP Parameters—Specifies the parameters needed for using an LDAP server. This area appears only when the selected server group uses LDAP.
  - Enable LDAP Over SSL—Specifies that SSL secures communications between the security appliance and the LDAP server. Also called secure LDAP.
  - Server Port—Specifies the server port to use. Enter the TCP port number by which you access the server.
  - Server Type—Lets you manually set the LDAP server type as a Sun Microsystems JAVA System Directory Server (formerly the Sun ONE Directory Server) or a Microsoft Active Directory, or lets you specify auto-detection for server type determination.
  - Base DN—Specifies the Base DN. Enter the location in the LDAP hierarchy where the server should begin searching when it receives an authorization request. For example, `OU=people, dc=cisco, dc=com`.
  - Scope—Specifies the extent of the search in the LDAP hierarchy that the server should make when it receives an authorization request—One Level (Search only one level beneath the Base DN. This option is quicker.) All Levels (Search all levels beneath the Base DN; in other words, search the entire subtree hierarchy. This option takes more time.)
  - Naming Attribute(s)—Specifies the Relative Distinguished Name attribute (or attributes) that uniquely identifies an entry on the LDAP server. Common naming attributes are Common Name (cn) and User ID (uid).
  - Login DN—Specifies the login DN. Some LDAP servers (including the Microsoft Active Directory server) require the security appliance to establish a handshake via authenticated binding before they will accept requests for any other LDAP operations. The security appliance identifies itself for authenticated binding by attaching a Login DN field to the user authentication request. The Login DN field defines the security appliance's authentication characteristics; these characteristics should correspond to those of a user with administration privileges. Enter the name of the directory object for security appliance authenticated binding, for example: cn=Administrator, cn=users, ou=people, dc=Example Corporation, dc=com. For anonymous access, leave this field blank.
  - Login Password—Specifies the login password. The characters you type are replaced with asterisks.
  - LDAP Attribute Map—Lists the LDAP attribute maps that you can apply to LDAP server. The LDAP attribute map translates Cisco attribute names into user-defined attribute names and values.
  - SASL MD5 authentication—Specifies that the MD5 mechanism of the Simple Authentication and Security Layer secures authentication communications between the security appliance and the LDAP server.
  - SASL Kerberos authentication—Specifies that Kerberos mechanism of the Simple Authentication and Security Layer secures authentication communications between the security appliance and the LDAP server.
  - Kerberos Server Group—Specifies the Kerberos server or server group used for authentication.
- NT Domain Parameters—Specifies the parameters needed for using an NT server and includes the following fields:

  - Server Port—Specifies the TCP port number by which you access the server. The default port number is 139.
  - NT Domain Controller— Specifies the NT Primary Domain Controller host name for this server, for example: PDC01. The maximum host name length is 15 characters. You must enter this name, and it must be the correct host name for the server for which you entered the IP Address in Authentication Server Address; if the name is incorrect, authentication fails.

- HTTP Form Parameters—Specifies the parameters for the HTTP Form protocol for single sign-on authentication, available only to WebVPN users. This area appears only when the selected server group uses HTTP Form, and only the Server Group name and the protocol are visible. Other fields are not available when using HTTP Form.

Note To configure SSO with the HTTP protocol correctly, you must have a thorough working knowledge of authentication and HTTP protocol exchanges.

If you do not know what the following parameters are, use an HTTP header analyzer to extract the data from the HTTP GET and POST exchanges when logging into the authenticating web server directly, not through the security appliance. See the WebVPN chapter in the *Cisco Security Appliance Command Line Configuration Guide* for more detail on extracting these parameters from the HTTP exchanges.

  - Start URL—Specifies the complete URL of the authenticating web server location where a pre-login cookie can be retrieved. This parameter must be configured only when the authenticating web server loads a pre-login cookie with the login page. A drop-down list offers both HTTP and HTTPS. The maximum number of characters is 1024, and there is no minimum.
  - Action URI—Specifies the complete Uniform Resource Identifier for the authentication program on the authorizing web server. The maximum number of characters for the complete URI is 2048 characters.
  - Username—Specifies the name of a username parameter—not a specific username—that must be submitted as part of the HTTP form used for SSO authentication. The maximum number of characters is 128, and there is no minimum.
  - Password—Specifies the name of a user password parameter—not a specific password value—that must be submitted as part of the HTTP form used for SSO authentication. The maximum number of characters is 128, and there is no minimum.
  - Hidden Values—Specifies hidden parameters for the HTTP POST request submitted to the authenticating web server for SSO authentication. This parameter is necessary only when it is expected by the authenticating web server as indicated by its presence in the HTTP POST request. The maximum number of characters is 2048.
  - Authentication Cookie Name—(Optional) Specifies the name of the cookie that is set by the server on successful login and that contains the authentication information. It is used to assign a meaningful name to the authentication cookie to help distinguish it from other cookies that the web server may pass back. The maximum number of characters is 128, and there is no minimum.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	• [1]	•	• [1.]	—

1. HTTP Form and WebVPN are supported only in single routed mode.

Test AAA Server

Configuration > Properties > AAA Setup > AAA Server Groups > Test

Note Test AAA Server is not available for HTTP Form authentication servers.

Use the Test button to determine whether the security appliance can contact the selected AAA server. Failure to reach the AAA server may be due to incorrect configuration in ASDM or the AAA server may be unreachable for other reasons, such as restrictive network configurations or server downtime.

After you complete the fields in this dialog box and click OK, the security appliance sends the applicable test message to the selected server. If the test fails, ASDM displays an error message about the type of error encountered. If the error message suggests a configuration error in ASDM, correct the configuration and try the test again.

Tip Checking for basic network connectivity to the AAA server may save you time in troubleshooting. To test basic connectivity, click **Tools > Ping.**

Fields

- AAA Server Group—*Display only.* Shows the AAA server group that the selected AAA server belongs to.
- Host —*Display only.* Shows the hostname of the AAA server you selected.
- Authorization—Specifies that ASDM tests authorizing a user with the selected AAA server. If the server type selected does not support authorization, this radio button is not available. For example, the security appliance cannot support authorization with Kerberos servers.
- Authentication—Specifies that ASDM tests authenticating a user with the selected AAA server. If the server type selected does not support authentication, this radio button is not available. For example, the security appliance cannot support authentication with LDAP servers.
- Username—Specifies the username you want to use to test the AAA server. Make sure the username exists on the AAA server; otherwise, the test will fail.
- Password—Specifies the password for the username you entered in the Username field. The Password field is available only for authentication tests. Make sure the password is correct for the username entered; otherwise, the authentication test will fail.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	• [1]	•	•	—

1. HTTP Form and WebVPN are supported only in single routed mode.

Auth. Prompt

Configuration > Properties > AAA Setup > Auth. Prompt

The Auth. Prompt pane lets you specify text to display to the user during the AAA authentication challenge process.You can specify the AAA challenge text for HTTP, FTP, and Telnet access through the security appliance when requiring user authentication from TACACS+ or RADIUS servers. This text is primarily for cosmetic purposes and displays above the username and password prompts that users view when logging in.

If the user authentication occurs from Telnet, you can use the User accepted message and User rejected message options to display different status prompts to indicate that the authentication attempt is accepted or rejected by the AAA server.

If the AAA server authenticates the user, the security appliance displays the User accepted message text, if specified, to the user; otherwise it displays the User rejected message text, if specified. Authentication of HTTP and FTP sessions displays only the challenge text at the prompt. The User accepted message and User rejected message text are not displayed.

Note Microsoft Internet Explorer displays up to 37 characters in an authentication prompt. Netscape Navigator displays up to 120 characters, and Telnet and FTP display up to 235 characters in an authentication prompt.

Fields

- Prompt—(Optional) Enables the display of AAA challenge text, specified in the field below the check box, for through-the-security appliance user sessions.
- Text—(Optional) Specify a string of up to 235 alphanumeric characters or 31 words, limited by whichever maximum is first reached. Do not use special characters; however, spaces and punctuation characters are permitted. Entering a question mark or pressing the Enter key ends the string. (The question mark appears in the string.)
- User accepted message—(Optional) Enables the display of text, specified in the field below the check box, confirming that the user has been authenticated.
- User rejected message—(Optional) Enables the display of text, specified in the field below the check box, indicating that authentication failed.

Note All of the fields in this pane are optional. If you do not specify an authentication prompt, FTP users see `FTP authentication`, HTTP users see `HTTP Authentication` Telnet users see no challenge text.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

LDAP Attribute Map

Configuration > Properties > AAA Setup > LDAP Attribute Map

The LDAP Attribute Map pane lets you create and name an attribute map for mapping custom (user-defined) attribute names to Cisco LDAP attribute names. If you are introducing a security appliance to an existing LDAP directory, your existing custom LDAP attribute names and values are probably different from the Cisco attribute names and values. Rather than renaming your existing attributes, you can create LDAP attribute maps that map your custom attribute names and values to Cisco attribute names and values. By using simple string substitution, the security appliance then presents you with only your own custom names and values.

You can then bind these attribute maps to LDAP servers or remove them as needed. You can also delete entire attribute maps or remove individual name and value entries.

Note To use the attribute mapping features correctly, you need to understand the Cisco LDAP attribute names and values as well as the user-defined attribute names and values.

Fields

- Name—Displays the names of the LDAP attribute maps available for editing.
- Attribute Map Name—Displays the mappings of custom attribute names to Cisco attribute names within each attribute map.
- Add—Displays the Add LDAP Attribute Map dialog box.
- Edit—Displays the Edit LDAP Attribute Map dialog box.
- Delete—Deletes the selected LDAP Attribute Map.

Add/Edit LDAP Attribute Map

Configuration > Properties > AAA Setup > LDAP Attribute Map > Add/Edit LDAP Attribute Map

The Add/Edit LDAP Attribute Map dialog box lets you modify or delete an existing LDAP attribute map, add a new LDAP attribute map, and populate attribute maps with attribute name and value mappings.

Your typical steps to add a new attribute map using the LDAP Attribute Map dialog box would be as follows:

1. Create a new, unpopulated attribute map.
2. Populate the attribute map with name mappings that translate Cisco attribute names to custom, user-defined attribute names.
3. Populate the attribute map with value mappings that apply custom, user-defined attribute values to the custom attribute name and to the matching Cisco attribute name and value.

You would then bind the attribute map to an LDAP server when adding or editing the LDAP server using the Add/Edit AAA Server dialog box.

Fields

- Name—Specifies the name of the LDAP attribute map you are adding or editing. If you are adding a new map, you enter the name of the map in this field. If you are editing a map that was selected in the LDAP Attribute Map pane, the name of the selected map displays as read-only text in this field. To change the map, you must return to the LDAP Attribute Map pane and choose the desired map.
- Name Map—Displays the fields necessary for mapping custom attribute names to Cisco attribute names.
- Value Map—Displays the fields necessary for mapping custom attribute values to custom attribute names and to the matching Cisco attribute name and value.

Add/Edit LDAP Attribute Map > Map Name Tab

Configuration > Properties > AAA Setup > LDAP Attribute Map > Add/Edit LDAP Attribute Map > Map Name Tab

The Add/Edit LDAP Attribute Map dialog box lets you modify or delete an existing LDAP attribute map, add a new LDAP attribute map, and populate attribute maps with attribute name and value mappings. See also Add/Edit LDAP Attribute Map.

Some fields vary depending upon whether you have selected the Map Name tab or the Map Value tab. When you click the Map Name tab, the following fields display.

Fields

- Name—Specifies the name of the LDAP attribute map you are adding or editing. If you are adding a new map, you enter the name of the map in this field. If you are editing a map that was selected in the LDAP Attribute Map pane, the name of the selected map displays as read-only text in this field. To change the map, you must return to the LDAP Attribute Map pane and choose the desired map.
- Custom Name—Specifies the custom, user-defined attribute name that maps to an attribute name selected from the Cisco Name drop-down list.
- Cisco Name—Specifies the Cisco attribute name you want to map to the user-defined name in the Custom Name field.
- Add—Inserts the name mapping into the attribute map.
- Remove—Removes the selected name mapping from the attribute map.
- Custom Name—Displays the custom attribute names of mappings in the attribute map.
- Cisco Name—Displays the Cisco attribute names of mappings in the attribute map.

Add/Edit LDAP Attribute Map > Map Value Tab

Configuration > Properties > AAA Setup > LDAP Attribute Map > Add/Edit LDAP Attribute Map > Map Value Tab

The Add/Edit LDAP Attribute Map dialog box lets you modify or delete an existing LDAP attribute map, add a new LDAP attribute map, and populate attribute maps with attribute name and value mappings. See also Add/Edit LDAP Attribute Map.

Some fields vary depending upon whether you have selected the Map Name tab or the Map Value tab. When you click the Map Value tab, the following fields appear.

Fields

- Name—Specifies the name of the LDAP attribute map you are adding or editing. If you are adding a new map, you enter the name of the map in this field. If you are editing a map that was selected in the LDAP Attribute Map pane, the name of the selected map displays as read-only text in this field. To change the map, you must return to the LDAP Attribute Map pane and choose the desired map.
- Custom Name—Displays the custom attribute names of mappings in the attribute map.
- Custom to Cisco Map Value—Displays the mapping of a custom value to a Cisco value for a custom attribute.
- Add—Displays the Add LDAP Attributes Map Value dialog box.
- Edit—Displays the Edit LDAP Attributes Map Value dialog box.
- Delete—Deletes the selected attribute value mapping from the LDAP attribute map.

Add/Edit LDAP Attributes Value Map

Configuration > Properties > AAA Setup > LDAP Attribute Map > Add/Edit LDAP Attribute Map > Add/Edit LDAP Attributes Map Value

The Add/Edit LDAP Attribute Map Value dialog box lets you map a custom attribute value for a custom attribute name to the Cisco value of the associated Cisco attribute name.

Fields

- Custom Name—If adding a new attribute value mapping, this is a drop-down list that lets you choose a custom attribute name from a list of attributes which do not yet have a custom value mapped to a Cisco attribute value. If editing an existing attribute value mapping, this is a read-only field which displays the name of the custom attribute selected on the Map Value tab of the Add/Edit LDAP Attribute Map dialog box.
- Custom Value—Specifies a custom value for the selected custom attribute.
- Cisco Value—Specifies the Cisco value for the selected custom attribute.
- Add—Adds the value mapping to the custom attribute value map.
- Remove—Removes the value mapping from the custom attribute value map.
- Custom Name—Displays the custom value for the custom attribute name.
- Cisco Name—Displays the Cisco value for the Cisco attribute name.

Configuring Device Access

Configuration > Device Access

This chapter contains the following topics:

AAA Access

Configuration > Device Access > AAA Access

The AAA Access pane includes tabs for configuring authentication, authorization, and accounting for management access. For an overview of AAA services, see Configuring AAA Servers.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Authentication Tab

Configuration > Device Access > AAA Access > Authentication Tab

Use this tab to enable authentication for administrator access to the security appliance. Authentication lets you control access by requiring a valid username and password. You can configure the security appliance to authenticate the following items:

- All administrative connections to the security appliance using the following methods:
  - Telnet
  - SSH
  - HTTPS/ASDM
  - Serial
- The **enable** command

Fields

- Require authentication to allow use of privileged mode commands—Specifies the parameters that control access to the privileged mode commands.
  - Enable—Enables or disables the requirement that a user be authenticated before being allowed to use privileged mode commands.
  - Server Group—Selects the server group to use for authenticating users to use privileged mode commands.
  - Use LOCAL when server group fails—Allows the use of the LOCAL database for authenticating users to use privileged mode commands if the selected server group fails.
- Require authentication for the following types of connections—Specifies the types of connections for which you want to require authentication and specifies the server group to use for that authentication.
  - HTTP/ASDM—Specifies whether to require authentication for HTTP/ASDM connections.
  - Server Group—Selects the server group to use for authenticating the specified connection type.
  - Use LOCAL when server group fails—Allows the use of the LOCAL database for authenticating the specified connection type if the selected server group fails.
  - SSH—Specifies whether to require authentication for SSH connections.
  - Telnet—Specifies whether to require authentication for Telnet connections.
  - Serial—Specifies whether to require authentication for serial connections.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Authorization Tab

Configuration > Device Access > AAA Access > Authorization Tab

Authorization lets you control access *per user* after you authenticate with a valid username and password. You can configure the security appliance to authorize management commands.

Authorization lets you control which services and commands are available to an individual user. Authentication alone provides the same access to services for all authenticated users.

When you enable command authorization, you have the option of manually assigning privilege levels to individual commands or groups of commands (using the **Advanced** button) or enabling the Predefined User Account Privileges (using the Restore Predefined User Account Privileges button):

Predefined User	Privilege Level	Description
Admin	15	Full access to all CLI commands
Read Only	5	Read only access to all commands
Monitor Only	3	Monitoring tab only

The Predefined User Account Privileges Setup pane displays a list of commands and privileges ASDM issues to the security appliance if you click Yes. Yes allows ASDM to support the three privilege levels: Admin, Read Only and Monitor Only.

The Command Privileges Setup pane displays a list of commands and privileges ASDM is going to issue to the security appliance. You can select one or more commands in the lists and use the Edit button to change the privilege level for the selected commands.

Fields

- Enable—Enables or disables authorization for security appliance command access. Selecting this check box activates the remaining parameters on this pane.
- Server Group—Selects the server group to use for authorizing users for command access.
- Use LOCAL when server group fails—Allows the use of the LOCAL database for authorizing users to use privileged mode commands if the selected server group fails.
- Advanced—Opens the Command Privileges Setup pane, on which you can manually assign privilege levels to individual commands or a group of commands.
- Restore Predefined User Account Privileges—Opens the Predefined User Account Command Privilege Setup pane, which sets up predefined user profiles and sets the privilege levels for the selected, listed commands.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Command Privileges Setup

Configuration > Device Access > AAA Access > Authorization > Command Privileges Setup

Use this pane to assign privilege levels to individual commands or groups of commands. Clicking a column heading sorts the entire table in alphanumeric order, using the selected column as the key field.

- Command Mode—Selects a specific command mode or All Modes. This selection determines what appears in the Command Modes table immediately below this list.
- CLI Command—Specifies the name of a CLI command.
- Mode—Indicates a mode that applies to this command. Certain commands have more than one mode.
- Variant—Indicates the form (for example, show or clear) of the specified command to which the privilege level applies.
- Privilege—Shows the privilege level currently assigned to this command.
- Edit—Displays the Select Command(s) Privilege dialog box. This dialog box lets you select from a list the privilege level for one or more commands selected on the parent window. The Command Modes table reflects the change as soon as you click OK.
- Select All—Selects the entire contents of the Command Modes table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Predefined User Account Command Privilege Setup

Configuration > Device Access > AAA Access > Authorization > Predefined User Account Command Privilege Setup

This pane asks whether you want the security appliance to set up user profiles named Admin, Read Only, and Monitor Only. You get to this pane by clicking Restore Predefined user Account Privileges on the Authorization tab of the Authentication/Authorization/Accounting pane.

Fields

- Command List—Lists the CLI commands, their modes, variants, and privileges, affected by the predefined user account privilege setup.
  - CLI Command—Specifies the name of a CLI command.
  - Mode—Indicates a mode that applies to this command. Certain commands have more than one mode.
  - Variant—Indicates the form (for example, show or clear) of the specified command to which the privilege level applies.
  - Privilege—Shows the privilege level currently assigned to this command.
- Yes—Directs the security appliance to set up the listed commands with the respective privilege levels. This setup lets you create users through the User Accounts pane with the roles Admin, privilege level 15; Read Only, with privilege level 5; and Monitor Only with privilege level 3.

- No—Lets you manage the privilege levels of commands and users manually.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Accounting Tab

Configuration > Device Access > AAA Access > Accounting Tab

Accounting lets you keep track of traffic that passes through the security appliance. If you enable authentication for that traffic, you can account for traffic per user. If you do not authenticate the traffic, you can account for traffic per IP address. Accounting information includes when sessions start and stop, the AAA client messages and username, the number of bytes that pass through the security appliance for the session, the service used, and the duration of each session.

Note You can configure accounting only for a TACACS+ server group. If no such group has yet been configured, go to Configuration > Properties > AAA Setup > AAA Server Groups.

Fields

- Require accounting to allow accounting of user activity—Specifies parameters related to accounting of user activity.
  - Enable—Enables or disables the requirement to allow accounting of user activity.
  - Server Group—Specifies the selected server group, if any, to use for user accounting. If no TACACS+ server group exists, the default value of this list is --None--.

Note The definition of the Server Group list parameter is the same for all group boxes on this pane.

- Require accounting for the following types of connections—Specifies the connection types for which you want to require accounting and the respective server groups for each.
  - HTTP/ASDM—Requires accounting for HTTP/ASDM connections.
  - Serial—Requires accounting for serial connections.
  - SSH—Requires accounting for secure shell (SSH) connections.
  - Telnet—Requires accounting for Telnet connections.
- Require command accounting for *Security Appliance*—You can send accounting messages to the TACACS+ accounting server when you enter any command other than **show** commands at the CLI. If you customize the command privilege level using the Configuration > Device Access > AAA Access > Authorization > Command Privilege Setup dialog box, you can limit which commands the security appliance accounts for by specifying a minimum privilege level. The security appliance does not account for commands that are below the minimum privilege level.

- Enable—Enables accounting for commands.
- Privilege level—Sets the minimum privilege level for which to perform command accounting.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

HTTPS/ASDM

Configure > Device Access > HTTPS/ASDM

The HTTPS/ASDM pane provides a table that specifies the addresses of all the hosts or networks that are allowed access to the ASDM using HTTPS. You can use this table to add or change the hosts or networks that are allowed access.

Fields

- Interface—Lists the interface on the security appliance from which the administrative access to the device manager is allowed.
- IP Address—Lists the IP address of the network or host that is allowed access.
- Mask—Lists the network mask associated with the network or host that is allowed access.
- Add—Displays the Add HTTP Configuration dialog box for adding a new host or network.
- Edit—Displays the Edit HTTP Configuration dialog box for editing the selected host or network.
- Delete—Deletes the selected host or network.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit HTTP Configuration

Configure > Device Access > HTTPS/ASDM > Add/Edit HTTP Configuration

The Add/Edit HTTP Configuration dialog box lets you add a host or network that will be allowed administrative access to the security appliance device manager over HTTPS.

Fields

- Interface Name—Specifies the interface on the security appliance from which the administrative access to the security appliance device manager is allowed.
- IP Address—Specifies the IP address of the network or host that is allowed access.
- Mask—Specifies the network mask associated with the network or host that is allowed access.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Secure Shell

Configuration > Device Access > HTTPS/ASDM > Secure Shell

The Secure Shell pane lets you configure rules that permit only specific hosts or networks to connect to the security appliance for administrative access using the SSH protocol. The rules restrict SSH access to a specific IP address and netmask. SSH connection attempts that comply with the rules must then be authenticated by a AAA server or the Telnet password.

You can monitor SSH sessions using Monitoring > Administration > Secure Shell Sessions.

Fields

The Secure Shell pane displays the following fields:

- Allowed SSH Versions—Restricts the version of SSH accepted by the security appliance. By default, SSH Version 1 and SSH Version 2 connections are accepted.
- Timeout (minutes)—Displays the number of minutes, 1 to 60, the Secure Shell session can remain idle before the security appliance closes it. The default is 5 minutes.
- SSH Access Rule—Displays the hosts and networks that are allowed to access the security appliance using SSH. Double-clicking a row in this table opens the Edit SSH Configuration dialog box for the selected entry.
  - Interface—Displays the name of a security appliance interface that will permit SSH connections.
  - IP Address—Displays the IP address of each host or network permitted to connect to this security appliance through the specified interface.
  - Mask—Displays the netmask for the IP address of each host or network permitted to connect to this security appliance through the specified interface.
- Add—Opens the Add SSH Configuration dialog box.
- Edit—Opens the Edit SSH Configuration dialog box.
- Delete—Deletes the selected SSH access rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit SSH Configuration

Configuration > Device Access > HTTPS/ASDM > Secure Shell > Add/Edit SSH Configuration

The Add SSH Configuration dialog box lets you add a new SSH access rule to the rule table. The Edit SSH Configuration dialog box lets you change an existing rule.

Fields

- Interface—Specifies the name of the security appliance interface that permits SSH connections.
- IP Address—Specifies the IP address of the host or network that is permitted to establish an SSH connection with the security appliance.
- Mask—The netmask of the host or network that is permitted to establish an SSH connection with the security appliance.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Telnet

Configuration > Device Access > Telnet

The Telnet pane lets you configure rules that permit only specific hosts or networks running ASDM to connect to the security appliance using the Telnet protocol.

The rules restrict administrative Telnet access through a security appliance interface to a specific IP address and netmask. Connection attempts that comply with the rules must then be authenticated by a preconfigured AAA server or the Telnet password. You can monitor Telnet sessions using Monitoring > Telnet Sessions.

Note Although a configuration file may contain more, there may be only five Telnet sessions active at the same time in single context mode. In multiple context mode, there may be only five Telnet sessions active per context.

Fields

The Telnet pane displays the following fields:

Telnet Rule Table:

- Interface—Displays the name of a security appliance interface which will permit Telnet connections, an interface on which is located a PC or workstation running ASDM.
- IP Address—Displays the IP address of each host or network permitted to connect to this security appliance through the specified interface.

Note This is not the IP address of the security appliance interface.

- Netmask—Displays the netmask for the IP address of each host or network permitted to connect to this security appliance through the specified interface.

Note This is not the IP address of the security appliance interface.

- Timeout—Displays the number of minutes, 1 to 60, the Telnet session can remain idle before the security appliance closes it. The default is 5 minutes.
- Add—Opens the Add Telnet Configuration dialog box.
- Edit—Opens the Edit Telnet Configuration dialog box.
- Delete—Deletes the selected item.
- Apply—Sends changes made in ASDM to the security appliance and applies them to the running configuration. Click Save to write a copy of the running configuration to Flash memory. Use the **File** menu to write a copy of the running configuration to Flash memory, a TFTP server, or a failover standby unit.
- Reset—Discards changes and reverts to the information displayed before changes were made or the last time you clicked **Refresh** or **Apply**. After Reset, use Refresh to make sure that information from the current running configuration is displayed.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Telnet Configuration

Configuration > Device Access > Telnet > Add/Edit Telnet Configuration

Adding Telnet Rules

To add a rule to the Telnet rule table, perform the following steps:

1. Click the **Add** button to open the **Telnet > Add** dialog box.

2. Click **Interface** to add a security appliance interface to the rule table.
3. In the IP Address box, enter the IP address of the host running ASDM which will be permitted Telnet access through this security appliance interface.

Note This is not the IP address of the security appliance interface.

4. In the Mask list, select or enter a netmask for the IP address to be permitted Telnet access.

Note This is not a mask for the IP address of the security appliance interface.

5. To return to the previous pane click:

- **OK**—Accepts changes and returns to the previous pane.
- **Cancel**—Discards changes and returns to the previous pane.
- **Help**—Provides more information.

Editing Telnet Rules

To edit a rule in the Telnet rule table, perform the following steps:

1. Click **Edit** to open the Telnet > Edit dialog box.
2. Click **Interface** to select a security appliance interface from the rule table.
3. In the IP Address field, enter the IP address of the host running ASDM which will be permitted Telnet access through this security appliance interface.

Note This is not the IP address of the security appliance interface.

4. In the Mask list, select or enter a netmask for the IP address to be permitted Telnet access.

Note This is not a mask for the IP address of the security appliance interface.

5. To return to the previous Window, click one of the following buttons:

- **OK**—Accepts changes and returns to the previous pane.
- **Cancel**—Discards changes and returns to the previous pane.
- **Help**—Provides more information.

Deleting Telnet Rules

To delete a rule from the Telnet table, perform the following steps:

1. Select a rule from the Telnet rule table.
2. Click **Delete**.

Applying Changes

Changes to the table made by Add, Edit, or Delete are not immediately applied to the running configuration. To apply or discard changes, click one of the following buttons:

1. **Apply**—Sends changes made in ASDM to the security appliance and applies them to the running configuration. Click **Save** to write a copy of the running configuration to Flash memory. Use the **File** menu to write a copy of the running configuration to Flash memory, a TFTP server, or a failover standby unit.
2. **Reset**—Discards changes and reverts to the information displayed before changes were made or the last time you clicked **Refresh** or **Apply**. After Reset, use **Refresh** to make sure that information from the current running configuration is displayed.

Fields

- Interface Name—Select the interface to allow Telnet access to the security appliance.
- IP Address—Enter the IP address of the host or network permitted to Telnet to the security appliance.
- Mask—Enter the subnet mask of the host or network permitted to Telnet to the security appliance.
- OK—Accepts changes and returns to the previous pane.
- Cancel—Discards changes and returns to the previous pane.
- Help—Provides more information.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Virtual Access

Configuration > Properties > Device Access > Virtual Access

The Virtual Access pane lets you configure a virtual Telnet server address on the security appliance to use for network access authentication.

Although you can configure network access authentication for any protocol or service, only HTTP, Telnet, or FTP provide an authentication challenge. A user must first authenticate with one of these services before other traffic requiring authentication is allowed through.

In some cases, you might not want to allow HTTP, Telnet, or FTP through the security appliance, but still need to authenticate other types of traffic. In those cases, you can create a virtual Telnet server on the security appliance. User connect to the security appliance using Telnet to the virtual Telnet IP address and the security appliance provides a Telnet prompt. When an unauthenticated user connects to the virtual Telnet IP address, the user is challenged for a username and password, and is then authenticated by the AAA server. Once authenticated, the user sees the message “Authentication Successful.” The user can then access other services that require authentication.

To log out from the security appliance, reconnect to the virtual IP address; you are prompted to log out.

Fields

- Virtual Telnet Server—Enter the virtual Telnet server IP address.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Failover

This section contains the following topics:

Understanding Failover

The Failover pane contains the settings for configuring failover on the security appliance. However, the Failover pane changes depending upon whether you are in multiple mode or single mode, and when you are in multiple mode, it changes based on the security context you are in.

Failover allows you to configure two security appliances so that one will take over operation if the other fails. Using a pair of security appliances, you can provide high availability with no operator intervention. The security appliance communicates failover information over a dedicated failover link. This failover link can be either a LAN-based connection or, on the PIX security appliance platform, a dedicated serial failover cable. The following information is communicated over the failover link:

- The failover state (active or standby).
- Hello messages (keep-alives).
- Network link status.
- MAC address exchange.
- Configuration replication.

Caution All information sent over the failover and Stateful Failover links is sent in clear text unless you secure the communication with a failover key. If the security appliance is used to terminate VPN tunnels, this information includes any usernames, passwords and preshared keys used for establishing the tunnels. Transmitting this sensitive data in clear text could pose a significant security risk. We recommend securing the failover communication with a failover key if you are using the security appliance to terminate VPN tunnels.

The security appliance supports two types of failover, Active/Standby and Active/Active. Additionally, failover can be stateful or stateless. For more information about the types of failover, see the following topics:

Active/Standby Failover

In an Active/Standby configuration, the active security appliance handles all network traffic passing through the failover pair. The standby security appliance does not handle network traffic until a failure occurs on the active security appliance. Whenever the configuration of the active security appliance changes, it sends configuration information over the failover link to the standby security appliance.

When a failover occurs, the standby security appliance becomes the active unit. It assumes the IP and MAC addresses of the previously active unit. Because the other devices on the network do not see any changes in the IP or MAC addresses, ARP entries do not change or time out anywhere on the network.

Active/Standby failover is available to security appliances in single mode or multiple mode.

Active/Active Failover

In an Active/Active failover configuration, both security appliances pass network traffic. Active/Active failover is only available to security appliances in multiple context mode.

To enable Active/Active failover on the security appliance, you need to create failover groups. If you enable failover without creating failover groups, you are enabling Active/Standby failover. A failover group is simply a logical group of one or more security contexts. You can create two failover groups on the security appliance. You should create the failover groups on the unit that will have failover group 1 in the active state. The admin context is always a member of failover group 1. Any unassigned security contexts are also members of failover group 1 by default.

As in Active/Standby failover, each unit in an Active/Active failover pair is given a primary or secondary designation. Unlike Active/Standby failover, this designation does not indicate which unit is active when both units start simultaneously. Each failover group in the configuration is given a primary or secondary role preference. This preference determines on which unit in the failover pair the contexts in the failover group appear in the active state when both units start simultaneously. You can have both failover groups be in the active state on a single unit in the pair, with the other unit containing the failover groups in the standby state. However, a more typical configuration is to assign each failover group a different role preference to make each one active on a different unit, balancing the traffic across the devices.

Initial configuration synchronization occurs when one or both units start. This synchronization occurs as follows:

- When both units start simultaneously, the configuration is synchronized from the primary unit to the secondary unit.
- When one unit starts while the other unit is already active, the unit that is starting up receives the configuration from the already active unit.

After both units are running, commands are replicated from one unit to the other as follows:

- Commands entered within a security context are replicated from the unit on which the security context appears in the active state to the peer unit.

Note A context is considered in the active state on a unit if the failover group to which it belongs is in the active state on that unit.

- Commands entered in the system execution space are replicated from the unit on which failover group 1 is in the active state to the unit on which failover group 1 is in the standby state.
- Commands entered in the admin context are replicated from the unit on which failover group 1 is in the active state to the unit on which failover group 1 is in the standby state.

Failure to enter the commands on the appropriate unit for command replication to occur will cause the configurations to be out of synchronization. Those changes may be lost the next time the initial configuration synchronization occurs.

In an Active/Active failover configuration, failover occurs on a failover group basis, not a system basis. For example, if you designate both failover groups as active on the primary unit, and failover group 1 fails, failover group 2 remains active on the primary unit, while failover group 1 becomes active on the secondary unit.

Note When configuring Active/Active failover, make sure that the combined traffic for both units is within the capacity of each unit.

Stateless (Regular) Failover

Stateless failover is also referred to as regular failover. In stateless failover, all active connections are dropped when a failover occurs. Clients need to reestablish connections when the new active unit takes over.

Stateful Failover

Note Stateful Failover is not supported on the ASA 5505 series adaptive security appliance.

When Stateful Failover is enabled, the active unit in the failover pair continually passes per-connection state information to the standby unit. After a failover occurs, the same connection information is available at the new active unit. Supported end-user applications are not required to reconnect to keep the same communication session.

Note The IP address and MAC address for the state and LAN failover links do not change at failover.

To use Stateful Failover, you must configure a state link to pass all state information to the standby unit. If you are using a LAN failover connection rather than the serial failover interface (available on the PIX security appliance platform only), you can use the same interface for the state link as the failover link. However, we recommend that you use a dedicated interface for passing state information the standby unit.

The following information is passed to the standby unit when Stateful Failover is enabled:

- NAT translation table.
- TCP connection table (except for HTTP), including the timeout connection.
- HTTP connection states (if HTTP replication is enabled).
- H.323, SIP, and MGCP UDP media connections.
- The system clock.

- The ISAKMP and IPSec SA table.

The following information is not copied to the standby unit when Stateful Failover is enabled:

- HTTP connection table (unless HTTP replication is enabled).
- The user authentication (uauth) table.
- The ARP table.
- Routing tables.

Configuring Failover with the High Availability and Scalability Wizard

The High Availability and Scalability Wizard steps you through the process of creating an Active/Active failover configuration, and Active/Standby failover configuration, or a VPN Cluster Load Balancing configuration.

See the following topics for information about using the High Availability and Scalability Wizard:

- Accessing and Using the High Availability and Scalability Wizard
- Configuring Active/Active Failover with the High Availability and Scalability Wizard
- Configuring Active/Standby Failover with the High Availability and Scalability Wizard
- Configuring VPN Load Balancing with the High Availability and Scalability Wizard
- Field Information for the High Availability and Scalability Wizard

Accessing and Using the High Availability and Scalability Wizard

To open the High Availability and Scalability Wizard, choose **Wizards > High Availability and Scalability Wizard** from the ASDM menu bar. The first screen of the wizard appears.

To move to the next screen of the wizard, click the **Next** button. You must complete the mandatory field of each screen before you can move to the next screen.

To move to a previous screen of the wizard, click the **Back** button. If information filled in on later screens of the wizard is not affected by the change you make to an earlier screen, that information remains on the screen as you move forward through the wizard again. You do not need to reenter it.

To leave the wizard at any time without saving any changes, click **Cancel**.

To send your configuration to the security appliance at the end of the wizard, click **Finish**.

Configuring Active/Active Failover with the High Availability and Scalability Wizard

The following procedure provides a high-level overview for configuring Active/Active failover using the High Availability and Scalability Wizard. Each step in the procedure corresponds with a wizard screen. Click **Next** after completing each step, except for the last step, before moving to the next step. Each step also contains a reference to additional information that you may need to complete the step.

Step 1 Choose **Configure Active/Active** failover on the Choose the type of failover configuration screen.

See Choose the Type of Failover Configuration for more information about this screen.

Step 2 Enter the IP address of the failover peer on the Check Failover Peer Connectivity and Compatibility screen. Click **Test Compatibility**. You will not be able to move to the next screen until all compatibility tests are passed.

See Check Failover Peer Connectivity and Compatibility for more information about this screen.

Step 3 If the security appliance or the failover peer are in single context mode, change them to multiple context mode on the Change Device to Multiple Mode screen. When you change the security appliance to multiple context mode, it will reboot. ASDM automatically reestablishes communication with the security appliance when it has finished rebooting.

See Change Device to Multiple Mode for more information about this screen.

Step 4 (PIX 500 series security appliance only) Select cable-based or LAN-based failover on the Select Failover Communication Media screen.

See Select Failover Communication Media for more information about this screen.

Step 5 Assign security contexts to failover groups on the Context Configuration screen. You can add and delete contexts on this screen.

See Security Context Configuration for more information about this screen.

Step 6 Define the Failover Link on the Failover Link Configuration screen.

See Failover Link Configuration for more information about this screen.

Step 7 (Not available on the ASA 5505 security appliance) Define the Stateful Failover link on the State Link Configuration screen.

See State Link Configuration for more information about this screen.

Step 8 Add standby addresses to the security appliance interfaces on the Standby Address Configuration screen.

See Standby Address Configuration for more information about this screen.

Step 9 Review your configuration on the Summary screen. If necessary, use the Back button to go to a previous screen and make changes.

See Summary for more information about this screen.

Step 10 Click **Finish**.

The failover configuration is sent to the security appliance and to the failover peer.

Configuring Active/Standby Failover with the High Availability and Scalability Wizard

The following procedure provides a high-level overview for configuring Active/Standby failover using the High Availability and Scalability Wizard. Each step in the procedure corresponds with a wizard screen. Click **Next** after completing each step, except for the last step, before moving to the next step. Each step also contains a reference to additional information that you may need to complete the step.

Step 1 Choose **Configure Active/Standby** failover on the Choose the type of failover configuration screen. Click next.

See Choose the Type of Failover Configuration for more information about this screen.

Step 2 Enter the IP address of the failover peer on the Check Failover Peer Connectivity and Compatibility screen. Click **Test Compatibility**. You will not be able to move to the next screen until all compatibility tests are passed.

See Check Failover Peer Connectivity and Compatibility for more information about this screen.

Step 3 (PIX 500 series security appliance only) Select cable-based or LAN-based failover on the Select Failover Communication Media screen.

See Select Failover Communication Media for more information about this screen.

Step 4 Define the Failover Link on the Failover Link Configuration screen.

See Failover Link Configuration for more information about this screen.

Step 5 (Not available on the ASA 5505 security appliance) Define the Stateful Failover link on the State Link Configuration screen.

See State Link Configuration for more information about this screen.

Step 6 Add standby addresses to the security appliance interfaces on the Standby Address Configuration screen.

See Standby Address Configuration for more information about this screen.

Step 7 Review your configuration on the Summary screen. If necessary, use the Back button to go to a previous screen and make changes.

See Summary for more information about this screen.

Step 8 Click **Finish**.

The failover configuration is sent to the security appliance and to the failover peer.

Configuring VPN Load Balancing with the High Availability and Scalability Wizard

The following procedure provides a high-level overview for configuring VPN cluster load balancing using the High Availability and Scalability Wizard. Each step in the procedure corresponds with a wizard screen. Click **Next** after completing each step, except for the last step, before moving to the next step. Each step also contains a reference to additional information that you may need to complete the step.

Step 1 Choose **Configure VPN Cluster Load Balancing** failover on the Choose the type of failover configuration screen.

See Choose the Type of Failover Configuration for more information about this screen.

Step 2 Configure the VPN load balancing settings on the VPN Cluster Load Balancing Configuration screen.

See VPN Cluster Load Balancing Configuration for more information about this screen.

Step 3 Review your configuration on the Summary screen. If necessary, use the Back button to go to a previous screen and make changes.

See Summary for more information about this screen.

Step 4 Click **Finish**.

The failover configuration is sent to the security appliance and to the failover peer.

Field Information for the High Availability and Scalability Wizard

The following dialogs are available in the High Availability and Scalability Wizard. You will not see every dialog box when you run through the wizard; each dialog box appears depending on the type of failover you are configuring and the hardware platform you are configuring it on.

Choose the Type of Failover Configuration

The Choose the Type of Failover Configuration screen lets you select the type of failover to configure.

Fields

The Choose the Type of Failover Configuration displays the following informational fields. These are useful for determining the failover capabilities of the security appliance.

- Hardware Model—(*Display only*) Displays the security appliance model number.
- No. of Interfaces—(*Display only*) Displays the number of interfaces available on the security appliance.
- No. of Modules—(*Display only*) Displays the number of modules installed on the security appliance.
- Software Version—(*Display only*) Displays the version of the platform software on the security appliance.
- Failover License—(*Display only*) Displays the type of failover license installed on the device. You may need to purchase an upgraded license to configure failover.
- Firewall Mode—(*Display only*) Displays the firewall mode (routed or transparent) and the context mode (single or multiple).

Choose the type of failover configuration you are configuring:

- Configure Active/Active Failover—Configures the security appliance for Active/Active failover.
- Configure Active/Standby Failover—Configures the security appliance for Active/Standby failover.
- Configure VPN Cluster Load Balancing—Configures the security appliance to participate in VPN load balancing as part of a cluster.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

Check Failover Peer Connectivity and Compatibility

The Check Failover Peer Connectivity and Compatibility screen lets you verify that the selected failover peer is reachable and compatible with the current unit. If any of the connectivity and compatibility tests fail, you must correct the problem before you can proceed with the wizard.

Fields

- Peer IP Address—Enter the IP address of the peer unit. This address does not have to be the failover link address, but it must be an interface that has ASDM access enabled on it.
- Test Compatibility—Click this button to perform the following connectivity and compatibility tests:
  - Connectivity test from this ASDM to the peer unit
  - Connectivity test from this firewall device to the peer firewall device
  - Hardware compatibility test
  - Software version compatibility
  - Failover license compatibility
  - Firewall mode compatibility

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

Change Device to Multiple Mode

The Change Device to Multiple Mode dialog box appears for Active/Active failover configuration only. Active/Active failover requires the security appliance to be in multiple context mode. This dialog box lets you convert a security appliance in single context mode to multiple context mode.

When you convert from single context mode to multiple context mode, the security appliance creates the system configuration and the admin context from the current running configuration.The admin context configuration is stored in the admin.cfg file. The conversion process does not save the previous startup configuration, so if the startup configuration differed from the running configuration, those differences are lost.

Converting the security appliance from single context mode to multiple context mode causes the security appliance to reboot. However the High Availability and Scalability Wizard restores connectivity with the newly created admin context and reports the status in the Devices Status field in this dialog box.

You need to convert both the current security appliance and the peer security appliance to multiple context mode before you can proceed.

Fields

- Change *device* To Multiple Context—Causes the security appliance to change to multiple context mode. *device* is the hostname of the security appliance.
- Change *device* (peer) To Multiple Context—Causes the peer unit to change to multiple context mode. *device* is the hostname of the security appliance.
- Device Status—(*Display only*) Displays the status of the security appliance while converting to multiple context mode.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

Select Failover Communication Media

The Select Failover Communication Media appears only on PIX 500 series security appliances. This screen lets you select between using a failover cable or LAN-based connection for the failover link.

Fields

- Use Failover Cable—Choose this option to use a dedicated failover cable for failover communication.
- Use LAN-based connection—Choose this option to use a network connection for failover communication.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

Security Context Configuration

The Security Context Configuration screen appears for Active/Active configuration only. The Security Context Configuration screen lets you assign security contexts to failover groups. It displays the security contexts currently configured on the device and lets you add new ones or remove existing ones as needed.

Although you can create security contexts on this screen, you cannot assign interfaces to those contexts or configure any other properties for them. To configure context properties and assign interfaces to a context, you need to use the System > Security Contexts pane.

Fields

- Name—Displays the name of the security context. To change the name, click the name and type a new name.
- Failover Group—Displays the failover group the context is assigned to. To change the failover group for a security context, click the failover group and select the new failover group number from the drop-down list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

Failover Link Configuration

The Failover Link Configuration screen only appears if you are configuring LAN-based failover; it does not appear if you are configuring a PIX 500 series security appliance for cable-based failover.

Fields

- LAN Interface—Choose the interface to use for failover communication from the drop-down list.
- Logical Name—Type a name for the interface.
- Active IP—Type the IP address used for the failover link on the unit that has failover group 1 in the active state.
- Standby IP—Type the IP address used for the failover link on the unit that has failover group 1 in the standby state.
- Subnet Mask—Type or select a subnet mask for the Active IP and Standby IP addresses.
- Secret Key—(Optional) Enter the key used to encrypt failover communication. If this field is left blank, failover communication, including any passwords or keys in the configuration sent during command replication, is in clear text.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

State Link Configuration

The State Link Configuration screen does not appear in the wizard for ASDM running on the ASA 5505 platform.

The State Link Configuration lets you enable Stateful Failover and configure the Stateful Failover link properties.

Fields

- Use the LAN link as the State Link—Choose this option to pass state information across the LAN-based failover link. This option is not available on PIX 500 series security appliances configured for cable-based failover.
- Disable Stateful Failover—Choose this option to disable Stateful Failover.
- Configure another interface for Stateful failover—Choose this option to configure an unused interface as the Stateful Failover interface.
  - State Interface—Choose the interface you want to use for Stateful Failover communication from the drop-down list.
  - Logical Name—Type the name for the Stateful Failover interface.
  - Active IP—Type the IP address for the Stateful Failover link on the unit that has failover group 1 in the active state.
  - Standby IP—Type the IP address for the Stateful Failover link on the unit that has failover group 1 in the standby state.
  - Subnet Mask—Type or select a subnet mask for the Active IP and Standby IP addresses.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

Standby Address Configuration

Use the Standby Address Configuration screen to assign standby addresses to the interface on the security appliance.

Fields

- Device/Interface—(Active/Standby failover) Displays the interfaces configured on the failover units. Click the plus sign (+) by a device name to displays the interfaces on that device. Click the minus sign (-) by a device name to hides the interfaces on that device.
- Device/Group/Context/Interface—(Active/Active failover) Displays the interfaces configured on the failover unit. The interfaces are grouped by context and the contexts are grouped by failover group. Click the plus sign (+) by a device, failover group, or context name to expand the list. Click the minus sign (-) by a device, failover group, or context name to collapse the list.

- Active IP—Double-click this field to edit or add an active IP address. Changes to this field also appear in the Standby IP field for the corresponding interface on the peer unit.
- Standby IP—Double-click this field to edit or add a standby IP address. Changes to this field also appear in the Active IP field for the corresponding interface on the peer unit.
- Is Monitored—Check this check box to enable health monitoring for that interface. Uncheck the check box to disable the health monitoring. By default, health monitoring of physical interfaces is enabled and health monitoring of virtual interfaces is disabled.
- ASR Group—Select the asynchronous group ID from the drop-down list. This setting is only available for physical interface. For virtual interfaces this field displays “None”.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

VPN Cluster Load Balancing Configuration

If you have a remote-client configuration in which you are using two or more security appliances connected to the same network to handle remote sessions, you can configure these devices to share their session load. This feature is called *load balancing*. Load balancing directs session traffic to the least loaded device, thus distributing the load among all devices. It makes efficient use of system resources and provides increased performance anodize availability.

Use the VPN Cluster Load Balancing Configuration screen to set parameters necessary for this device to participate in a load balancing cluster.

Note VPN load balancing runs only on security appliance models ASA 5520 and higher.

VPN load balancing requires an active 3DES/AES license. The security appliance checks for the existence of this crypto license before enabling load balancing. If it does not detect an active 3DES or AES license, the security appliance prevents the enabling of load balancing and also prevents internal configuration of 3DES by the load balancing system unless the license permits this usage.

Enabling load balancing involves:

- Configuring the load-balancing cluster by establishing a common virtual cluster IP address, UDP port (if necessary), and IPSec shared secret for the cluster. These values are identical for every device in the cluster.
- Configuring the participating device by enabling load balancing on the device and defining device-specific properties. These values vary from device to device.

Note Load balancing is effective only on remote sessions initiated with the Cisco VPN Client (Release 3.0 and later), the Cisco VPN 3002 Hardware Client (Release 3.5 and later), or the ASA 5505 operating as an Easy VPN Client. All other clients, including LAN-to-LAN connections, can connect to a security appliance on which load balancing is enabled, but the cannot participate in load balancing.

To implement load balancing, you group together logically two or more devices on the same private LAN-to-LAN network into a *virtual cluster*.

Fields

- **VPN Load Balancing**—Configures virtual cluster device parameters.
  - **Cluster IP Address**—Specifies the single IP address that represents the entire virtual cluster. Choose an IP address that is within the public subnet address range shared by all the security appliances in the virtual cluster.
  - **Cluster UDP Port**—Specifies the UDP port for the virtual cluster in which this device is participating. The default value is 9023. If another application is using this port, enter the UDP destination port number you want to use for load balancing.
  - **Enable IPSec Encryption**—Enables or disables IPSec encryption. If you select this check box, you must also specify and verify a shared secret.The security appliances in the virtual cluster communicate via LAN-to-LAN tunnels using IPSec. To ensure that all load-balancing information communicated between the devices is encrypted, select this check box.

Note When using encryption, you must have previously configured the load-balancing inside interface. If that interface is not enabled on the load-balancing inside interface, you get an error message when you try to configure cluster encryption.

If the load-balancing inside interface is enabled when you configured cluster encryption, but is disabled before you configure the participation of the device in the virtual cluster, you get an error message when you select the Participate in Load Balancing Cluster check box, and encryption is not enabled for the cluster.

  - **Shared Secret Key**—Specifies the shared secret to between IPSec peers when you enable IPSec encryption. The value you enter in the box appears as consecutive asterisk characters.
  - **Priority Of This Device**—Specifies the priority assigned to this device within the cluster. The range is from 1 to 10. The priority indicates the likelihood of this device becoming the virtual cluster master, either at start-up or when an existing master fails. The higher you set the priority (for example, 10), the more likely this device becomes the virtual cluster master.
  - **Public Interface Of This Device**—Specifies the name or IP address of the public interface for this device.
  - **Private Interface Of This Device**—Specifies the name or IP address of the private interface for this device.

Note If the devices in the virtual cluster are powered up at different times, the first device to be powered up assumes the role of virtual cluster master. Because every virtual cluster requires a master, each device in the virtual cluster checks when it is powered-up to ensure that the cluster has a virtual master. If none exists, that device takes on the role. Devices powered up and added to the cluster later become secondary devices. If all the devices in the virtual cluster are powered up simultaneously, the device with the highest priority setting becomes the virtual cluster master. If two or more devices in the virtual cluster are powered up simultaneously, and both have the highest priority setting, the one with the lowest IP address becomes the virtual cluster master.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Summary

The Summary screen displays the results of the configuration steps you performed in the previous wizard panels.

Fields

The configuration appears in the center of the screen. Verify your settings and click **Finish** to send your configuration to the device. If you are configuring failover, the configuration is also sent to the failover peer. If you need to change a setting, click **Back** until you reach the screen where you need to make the change. Make the change and click **Next** until you return to the Summary screen.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	•

Field Information for the Failover Panes

What displays on the failover pane depends upon the mode you are in (single or multiple context mode) and whether you are in the system execution space or in a security context.

This section contains the following topics:

- Failover-Multiple Mode, System

Failover - Single Mode

Configuration > Properties > Failover

The Failover pane contains the tabs where you can configure Active/Standby failover in single context mode. For more information about failover, see Understanding Failover. For more information about configuring the settings on each tab of the Failover pane, see the following information. Note that the Interfaces tabs changes based on whether you are in routed firewall mode or transparent firewall mode.

- Failover: Setup
- Failover: Interfaces (Routed Firewall Mode)
- Failover: Interfaces (Transparent Firewall Mode)
- Failover: Criteria
- Failover: MAC Addresses

Failover: Setup

Configuration > Properties > Failover > Setup

Use this tab to enable failover on the security appliance. You also designate the failover link and the state link, if using Stateful Failover, on this tab.

For more information about configuring failover in general, see Understanding Failover.

Fields

- Enable Failover—Checking this check box enables failover and lets you configure a standby security appliance.

Note The speed and duplex settings for the failover interface cannot be changed when Failover is enabled. To change these settings for the failover interface, you must configure them in the Configuration > Interfaces pane before enabling failover.

ASDM displays a dialog box asking if you want to configure the peer unit when you enable failover. This dialog box also appears when the Preferred Role setting or, on the PIX security appliance platform, the Enable LAN rather than serial cable failover setting changes.

  - Peer IP Address—Enter an IP address on the peer unit that ASDM can connect to. This field appears on the Do you want to configure the failover peer firewall dialog box.

- Use 32 hexadecimal character key—Check this check box to enter a hexadecimal value for the encryption key in the Shared Key box. Uncheck this check box to enter an alphanumeric shared secret in the Shared Key box.
- Shared Key—Specifies the failover shared secret or key for encrypted and authenticated communications between failover pairs.

  If you checked the Use 32 hexadecimal character key check box, then enter a hexadecimal encryption key. The key must be 32 hexadecimal characters (0-9, a-f).

If you unchecked the Use 32 hexadecimal character key check box, then enter an alphanumeric shared secret. The shared secret can be from 1 to 63 characters. Valid character are any combination of numbers, letters, or punctuation. The shared secret is used to generate the encryption key.

- Enable LAN rather than serial cable failover—(PIX security appliance platform only) Check this check box to enable LAN Failover. Uncheck this check box to use the dedicated serial cable as the failover link.
- LAN Failover—Contains the fields for configuring LAN Failover.
  - Interface—Specifies the interface used for failover communication. Failover requires a dedicated interface, however you can share the interface with Stateful Failover.

    Only unconfigured interfaces or subinterfaces are displayed in this list and can be selected as the LAN Failover interface. Once you specify an interface as the LAN Failover interface, you cannot edit that interface in the Configuration > Interfaces pane.
  - Active IP—Specifies the IP address for the failover interface on the active unit.
  - Subnet Mask—Specifies the mask for the failover interface on the primary and secondary unit.
  - Logical Name—Specifies the logical name of the interface used for failover communication.
  - Standby IP—Specifies the IP address used by the secondary unit to communicate with the primary unit
  - Preferred Role—Specifies whether the preferred role for this security appliance is as the primary or secondary unit in a LAN failover.
- State Failover—Contains the fields for configuring Stateful Failover.

Note Stateful Failover is not available on the ASA 5505 platform. This area does not appear on ASDM running on an ASA 5505 security appliance.

  - Interface—Specifies the interface used for state communication. You can choose an unconfigured interface or subinterface, the LAN Failover interface, or the Use Named option.

Note We recommend that you use two separate, dedicated interfaces for the LAN Failover interface and the Stateful Failover interface.

    If you choose an unconfigured interface or subinterface, you must supply the Active IP, Subnet Mask, Standby IP, and Logical Name for the interface.

    If you choose the LAN Failover interface, you do not need to specify the Active IP, Subnet Mask, Logical Name, and Standby IP values; the values specified for the LAN Failover interface are used.

    If you choose the Use Named option, the Logical Name field becomes a drop-down list of named interfaces. Choose the interface from this list. The Active IP, Subnet Mask, and Standby IP values do not need to be specified. The values specified for the interface are used. Be sure to specify a standby IP address for the selected interface on the Interfaces tab.

Note Because Stateful Failover can generate a large amount of traffic, performance for both Stateful Failover and regular traffic can suffer when you use a named interface.

- Active IP—Specifies the IP address for the Stateful Failover interface on the primary unit. This field is dimmed if the LAN Failover interface or Use Named option is selected in the Interface drop-down list.
- Subnet Mask—Specifies the mask for the Stateful Failover interfaces on the primary and secondary units. This field is dimmed if the LAN Failover interface or Use Named option is selected in the Interface drop-down list.
- Logical Name—Specifies the logical interface used for failover communication. If you selected the Use Named option in the Interface drop-down list, this field displays a list of named interfaces. This field is dimmed if the LAN Failover interface is selected in the Interface drop-down list.
- Standby IP—Specifies the IP address used by the secondary unit to communicate with the primary unit. This field is dimmed if the LAN Failover interface or Use Named option is selected in the Interface drop-down list.
- Enable HTTP replication—Selecting this check box enables Stateful Failover to copy active HTTP sessions to the standby firewall. If you do not allow HTTP replication, then HTTP connections are disconnected at failover. Disabling HTTP replication reduces the amount of traffic on the state link.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

For More Information

For more information about failover in general, see Understanding Failover.

Failover: Interfaces (Routed Firewall Mode)

Configuration > Properties > Failover > Interfaces

Use this tab to define the standby IP address for each interface on the security appliance and to specify whether the status of the interface should be monitored.

For more information about configuring failover in general, see Understanding Failover.

Fields

- Interface—Lists the interfaces on the security appliance and identifies their active IP address, standby IP address, and monitoring status.
  - Interface Name column—Identifies the interface name.
  - Active IP column—Identifies the active IP address for this interface.
  - Standby IP column—Identifies the IP address of the corresponding interface on the standby failover unit.
  - Is Monitored column—Specifies whether this interface is monitored for failure.

- Edit—Displays the Edit Failover Interface Configuration (Routed Firewall Mode) dialog box for the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

For More Information

For more information about failover in general, see Understanding Failover.

Edit Failover Interface Configuration (Routed Firewall Mode)

Configuration > Properties > Failover > Interfaces > Edit Failover Interface Configuration

Use the Edit Failover Interface Configuration dialog box to define the standby IP address for an interface and to specify whether the status of the interface should be monitored.

Fields

- Interface Name—Identifies the interface name.
- Active IP Address—Identifies the IP address for this interface. This field does not appear if an IP address has not been assigned to the interface.
- Subnet Mask—Identifies the mask for this interface. This field does not appear if an IP address has not been assigned to the interface.
- Standby IP Address—Specifies the IP address of the corresponding interface on the standby failover unit. This field does not appear if an IP address has not been assigned to the interface.
- Monitor interface for failure—Specifies whether this interface is monitored for failure. The number of interfaces that can be monitored for the security appliance is 250. Monitored failover interfaces can have the following status:
  - Unknown—Initial status. This status can also mean the status cannot be determined.
  - Normal—The interface is receiving traffic.
  - Testing—Hello messages are not heard on the interface for five poll times.
  - Link Down—The interface is administratively down.
  - No Link—The physical link for the interface is down.
  - Failed—No traffic is received on the interface, yet traffic is heard on the peer interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

For More Information

For more information about failover in general, see Understanding Failover.

Failover: Interfaces (Transparent Firewall Mode)

Configuration > Properties > Failover > Interfaces

Use this tab to define the standby management IP address and to specify whether the status of the interfaces on the security appliance should be monitored.

Fields

- Interface—Lists the interfaces on the security appliance and identifies their monitoring status.
  - Interface Name column—Identifies the interface name.
  - Is Monitored column—Specifies whether this interface is monitored for failure.
- Edit—Displays the Edit Failover Interface Configuration (Transparent Firewall Mode) dialog box for the selected interface.
- Management IP Address—Identifies the active and standby management IP addresses for the security appliance or for a context in transparent firewall mode.
  - Active—Identifies the active management IP address.
  - Standby—Specifies the management IP address on the standby failover unit.
- Management Netmask—Identifies the mask associated with the active and standby management IP addresses.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
—	•	•	—	—

For More Information

For more information about failover in general, see Understanding Failover.

Edit Failover Interface Configuration (Transparent Firewall Mode)

Configuration > Properties > Failover > Interfaces > Edit Failover Interface Configuration

Use the Edit Failover Interface Configuration dialog box to specify whether the status of the interface should be monitored.

Fields

- Interface Name—Identifies the interface name.
- Monitor interface for failure—Specifies whether this interface is monitored for failure. The number of interfaces that can be monitored for the security appliance is 250. Hello messages are exchanged between the security appliance failover pair during every interface poll time period. Monitored failover interfaces can have the following status:
  - Unknown—Initial status. This status can also mean the status cannot be determined.
  - Normal—The interface is receiving traffic.
  - Testing—Hello messages are not heard on the interface for five poll times.
  - Link Down—The interface is administratively down.
  - No Link—The physical link for the interface is down.
  - Failed—No traffic is received on the interface, yet traffic is heard on the peer interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
—	•	•	—	—

For More Information

For more information about failover in general, see Understanding Failover.

Failover: Criteria

Configuration > Properties > Failover > Criteria

Use this tab to define criteria for failover, such as how many interfaces must fail and how long to wait between polls. The hold time specifies the interval to wait without receiving a response to a poll before unit failover.

Fields

- Interface Policy—Contains the fields for defining the policy for failover when monitoring detects an interface failure.
  - Number of failed interfaces that triggers failover—When the number of failed monitored interfaces exceeds the value you set with this command, then the security appliance fails over. The range is between 1 and 250 failures.
  - Percentage of failed interfaces that triggers failover—When the number of failed monitored interfaces exceeds the percentage you set with this command, then the security appliance fails over.

- Failover Poll Times—Contains the fields for defining how often hello messages are sent on the failover link, and, optionally, how long to wait before testing the peer for failure if no hello messages are received.
  - Unit Failover—The amount of time between hello messages among units. The range is between 1 and 15 seconds or between 200 and 999 milliseconds.
  - Unit Hold Time—Sets the time during which a unit must receive a hello message on the failover link, or else the unit begins the testing process for peer failure. The range is between 1and 45 seconds or between 800 and 999 milliseconds. You cannot enter a value that is less than 3 times the polltime.
  - Monitored Interfaces—The amount of time between polls among interfaces. The range is between 1and 15 seconds or 500 to 999 milliseconds.
  - Interface Hold Time—Sets the time during which a data interface must receive a hello message on the data interface, after which the peer is declared failed. Valid values are from 5 to 75 seconds.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

For More Information

For more information about failover in general, see Understanding Failover.

Failover: MAC Addresses

Configuration > Properties > Failover > MAC Addresses

The MAC Addresses tab lets you configure the virtual MAC addresses for the interfaces in an Active/Standby failover pair.

Note This tab is not available on the ASA 5505 platform.

In Active/Standby failover, the MAC addresses for the primary unit are always associated with the active IP addresses. If the secondary unit boots first and becomes active, it uses the burned-in MAC address for its interfaces. When the primary unit comes online, the secondary unit obtains the MAC addresses from the primary unit. The change can disrupt network traffic.

You can configure virtual MAC addresses for each interface to ensure that the secondary unit uses the correct MAC addresses when it is the active unit, even if it comes online before the primary unit. If you do not specify virtual MAC addresses, then the failover pair uses the burned-in NIC address as the MAC address.

Note You cannot configure a virtual MAC address for the failover or state links. The MAC and IP addresses for those links do not change during failover.

Fields

- MAC Addresses—Lists physical interfaces on the security appliance for which an active and standby virtual MAC address has been configured.
  - Physical Interface column—Identifies the physical interface for which failover virtual MAC addresses are configured.
  - Active MAC Address column—Identifies the MAC address of the active security appliance (usually primary).
  - Standby MAC Address column—Identifies the MAC address of the standby security appliance (usually secondary).
- Add—Displays the Add Interface MAC Address dialog box. You cannot assign virtual MAC addresses to the LAN failover and Stateful Failover interfaces. See Add/Edit Interface MAC Address for more information.
- Edit—Displays the Edit Interface MAC Address dialog box for the selected interface. See Add/Edit Interface MAC Address for more information.
- Delete—Removes the currently selected interface from the MAC addresses table. There is no confirmation or undo.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

For More Information

For more information about failover in general, see Understanding Failover.

Add/Edit Interface MAC Address

Configuration > Properties > Failover > MAC Addresses > Add/Edit Interface MAC Address

Use the Add/Edit Interface MAC Address dialog box to define the active and standby virtual MAC addresses for an interface.

Fields

- Physical Interface—Specifies the physical interface for which you are defining failover virtual MAC addresses. Because the MAC addresses do not change for the LAN failover and Stateful Failover interfaces during failover, you cannot choose these interfaces.
- MAC Addresses—Contains the fields for specifying the active and standby virtual MAC addresses for the interface.
  - Active Interface—Specifies the MAC address of the interface on the active security appliance (usually primary). Enter the MAC address in hexadecimal format (for example, 0123.4567.89AB).

– Standby Interface—Specifies the MAC address of the interface on the standby security appliance (usually secondary). Enter the MAC address in hexadecimal format (for example, 0123.4567.89AB).

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

For more information about failover in general, see Understanding Failover.

Failover-Multiple Mode, Security Context

The fields displayed on the Failover pane in multiple context mode change depending upon whether the context is in transparent or routed firewall mode.

This section contains the following topics:

- Failover - Routed
- Failover - Transparent

Failover - Routed

Configuration > Properties > Failover

Use this pane to define the standby IP address for each interface in the security context and to specify whether the status of the interface should be monitored.

Fields

- Interface table—Lists the interfaces on the security appliance and identifies their active IP address, standby IP address, and monitoring status.
  - Interface Name column—Identifies the interface name.
  - Active IP column—Identifies the active IP address for this interface.
  - Standby IP column—Identifies the IP address of the corresponding interface on the standby failover unit.
  - Is Monitored column—Specifies whether this interface is monitored for failure.
- Edit—Displays the Edit Failover Interface Configuration dialog box for the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	—	•	—

For More Information

For more information about failover in general, see Understanding Failover.

Edit Failover Interface Configuration

Configuration > Properties > Failover > Edit Failover Interface Configuration

Use the Edit Failover Interface Configuration dialog box to define the standby IP address for an interface and to specify whether the status of the interface should be monitored.

Fields

- Interface Name—Identifies the interface name.
- Active IP Address—Identifies the IP address for this interface. This field does not appear if an IP address has not been assigned to the interface.
- Subnet Mask—Identifies the mask for this interface. This field does not appear if an IP address has not been assigned to the interface.
- Standby IP Address—Specifies the IP address of the corresponding interface on the standby failover unit. This field does not appear if an IP address has not been assigned to the interface.
- Monitor interface for failure—Specifies whether this interface is monitored for failure. The number of interfaces that can be monitored for the security appliance is 250. Hello messages are exchanged between the security appliance failover pair during every interface poll time period. Monitored failover interfaces can have the following status:
  - Unknown—Initial status. This status can also mean the status cannot be determined.
  - Normal—The interface is receiving traffic.
  - Testing—Hello messages are not heard on the interface for five poll times.
  - Link Down—The interface is administratively down.
  - No Link—The physical link for the interface is down.
  - Failed—No traffic is received on the interface, yet traffic is heard on the peer interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	—	•	—

For More Information

For more information about failover in general, see Understanding Failover.

Failover - Transparent

Configuration > Properties > Failover

Use this pane to define the standby IP address for the management interface for the security context and to specify whether the status of the interfaces on the security context should be monitored.

Fields

- Interface—Lists the interfaces for the security context and identifies their monitoring status.
  - Interface Name—Identifies the interface name.
  - Is Monitored—Specifies whether this interface is monitored for failure.
- Edit—Displays the Edit Failover Interface Configuration dialog box for the selected interface.
- Management IP Address—Identifies the active and standby management IP addresses for the security context.
  - Active—Identifies the management IP address for the active failover unit.
  - Standby—Specifies the management IP address for the standby failover unit.
- Management Netmask—Identifies the mask associated with the management address.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
—	•	—	•	—

For More Information

For more information about failover in general, see Understanding Failover.

Edit Failover Interface Configuration

Configuration > Properties > Failover > Interfaces > Edit Failover Interface Configuration

Use the Edit Failover Interface Configuration dialog box to specify whether the status of the interface should be monitored.

Fields

- Interface Name—Identifies the interface name.
- Monitor interface for failure—Specifies whether this interface is monitored for failure. The number of interfaces that can be monitored for the security appliance is 250. Hello messages are exchanged between the security appliance failover pair during every interface poll time period. Monitored failover interfaces can have the following status:
  - Unknown—Initial status. This status can also mean the status cannot be determined.

- Normal—The interface is receiving traffic.
- Testing—Hello messages are not heard on the interface for five poll times.
- Link Down—The interface is administratively down.
- No Link—The physical link for the interface is down.
- Failed—No traffic is received on the interface, yet traffic is heard on the peer interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
—	•	—	•	—

For More Information

For more information about failover in general, see Understanding Failover.

Failover-Multiple Mode, System

System > Configuration > Failover

This pane includes tabs for configuring the system-level failover settings in the system context of a security appliance in multiple context mode. In multiple mode, you can configure Active/Standby or Active/Active failover. Active/Active failover is automatically enabled when you create failover groups in the device manager. For both types of failover, you need to provide system-level failover settings in the system context, and context-level failover settings in the individual security contexts. For more information about configuring failover in general, see Understanding Failover.

Seethe following topics for more information:

- Failover > Setup Tab
- Failover > Criteria Tab
- Failover > Active/Active Tab
- Failover > MAC Addresses Tab

Failover > Setup Tab

System > Configuration > Failover > Setup Tab

Use this tab to enable failover on a security appliance in multiple context mode. You also designate the failover link and the state link, if using Stateful Failover, on this tab.

Fields

- Enable Failover—Checking this check box enables failover and lets you configure a standby security appliance.

Note The speed and duplex settings for an interface cannot be changed when Failover is enabled. To change these settings for the failover interface, you must configure them in the Configuration > Interfaces pane before enabling failover.

- Use 32 hexadecimal character key—Check this check box to enter a hexadecimal value for the encryption key in the Shared Key field. Uncheck this check box to enter an alphanumeric shared secret in the Shared Key field.
- Shared Key—Specifies the failover shared secret or key for encrypted and authenticated communications between failover pairs.

  If you checked the Use 32 hexadecimal character key check box, then enter a hexadecimal encryption key. The key must be 32 hexadecimal characters (0-9, a-f).

  If you cleared the Use 32 hexadecimal character key check box, then enter an alphanumeric shared secret. The shared secret can be from 1 to 63 characters. Valid character are any combination of numbers, letters, or punctuation. The shared secret is used to generate the encryption key.
- Enable LAN rather than serial cable failover—(PIX security appliance platform only) Check this check box to enable LAN failover. Uncheck this check box to use the dedicated serial link as the failover link.
- LAN Failover—Contains the fields for configuring LAN Failover.
  - Interface—Specifies the interface used for failover communication. Failover requires a dedicated interface, however, you can use the same interface for Stateful Failover.

    Only unconfigured interfaces or subinterfaces that have not been assigned to a context are displayed in this list and can be selected as the LAN Failover interface. Once you specify an interface as the LAN Failover interface, you cannot edit that interface in the Configuration > Interfaces pane or assign that interface to a context.
  - Active IP—Specifies the IP address for the failover interface on the active unit.
  - Subnet Mask—Specifies the mask for the failover interface on the active unit.
  - Logical Name—Specifies the logical name for the failover interface.
  - Standby IP—Specifies the IP address of the standby unit.
  - Preferred Role—Specifies whether the preferred role for this security appliance is as the primary or secondary unit in a LAN failover.
- State Failover—Contains the fields for configuring Stateful Failover.
  - Interface—Specifies the interface used for failover communication. You can choose an unconfigured interface or subinterface or the LAN Failover interface.

    If you choose the LAN Failover interface, the interface needs enough capacity to handle both the LAN Failover and Stateful Failover traffic. Also, you do not need to specify the Active IP, Subnet Mask, Logical Name, and Standby IP values; the values specified for the LAN Failover interface are used.

    Note We recommend that you use two separate, dedicated interfaces for the LAN Failover interface and the Stateful Failover interface.

  - Active IP—Specifies the IP address for the Stateful Failover interface on the active unit.
  - Subnet Mask—Specifies the mask for the Stateful Failover interface on the active unit.

- Logical Name—Specifies the logical name for the Stateful Failover interface.
- Standby IP—Specifies the IP address of the standby unit.
- Enable HTTP replication—Checking this check box enables Stateful Failover to copy active HTTP sessions to the standby firewall. If you do not allow HTTP replication, then HTTP connections are disconnected at failover. Disabling HTTP replication reduces the amount of traffic on the state link.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

For More Information

For more information about failover in general, see Understanding Failover.

Failover > Criteria Tab

System > Configuration > Failover > Criteria Tab

Use this tab to define criteria for failover, such as how many interfaces must fail and how long to wait between polls. The hold time specifies the interval to wait without receiving a response to a poll before unit failover.

Note If you are configuring Active/Active failover, you do not use this tab to define the interface policy; instead, you define the interface policy for each failover group using the Failover > Active/Active Tab. With Active/Active failover, the interface policy settings defined for each failover group override the settings on this tab. If you disable Active/Active failover, then the settings on this tab are used.

Fields

- Interface Policy—Contains the fields for defining the policy for failover when monitoring detects an interface failure.
  - Number of failed interfaces that triggers failover—When the number of failed monitored interfaces exceeds the value you set with this command, then the security appliance fails over. The range is between 1 and 250 failures.
  - Percentage of failed interfaces that triggers failover—When the number of failed monitored interfaces exceeds the percentage you set with this command, then the security appliance fails over.
- Failover Poll Times—Contains the fields for defining how often hello messages are sent on the failover link, and, optionally, how long to wait before testing the peer for failure if no hello messages are received.
  - Unit Failover—The amount of time between hello messages among units. The range is between 1 and 15 seconds or between 200 and 999 milliseconds.

- Unit Hold Time—Sets the time during which a unit must receive a hello message on the failover link, or else the unit begins the testing process for peer failure. The range is between 1and 45 seconds or between 800 and 999 milliseconds. You cannot enter a value that is less than 3 times the polltime.
- Monitored Interfaces—The amount of time between polls among interfaces. The range is between 1and 15 seconds or 500 to 999 milliseconds.
- Interface Hold Time—Sets the time during which a data interface must receive a hello message on the data interface, after which the peer is declared failed. Valid values are from 5 to 75 seconds.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

For More Information

For more information about failover in general, see Understanding Failover.

Failover > Active/Active Tab

System > Configuration > Failover > Active/Active Tab

Use this tab to enable Active/Active failover on the security appliance by defining failover groups. In an Active/Active failover configuration, both security appliances pass network traffic. Active/Active failover is only available to security appliances in multiple mode.

A failover group is simply a logical group of security contexts. You can create two failover groups on the security appliance. You must create the failover groups on the active unit in the failover pair. The admin context is always a member of failover group 1. Any unassigned security contexts are also members of failover group 1 by default.

Note When configuring Active/Active failover, make sure that the combined traffic for both units is within the capacity of each unit.

Fields

- Failover Groups—Lists the failover groups currently defined on the security appliance.
  - Group Number—Specifies the failover group number. This number is used when assigning contexts to failover groups.
  - Preferred Role—Specifies the unit in the failover pair, primary or secondary, on which the failover group appears in the active state when both units start up simultaneously or when the preempt option is specified. You can have both failover groups be in the active state on a single unit in the pair, with the other unit containing the failover groups in the standby state. However, a more typical configuration is to assign each failover group a different role preference to make each one active on a different unit, balancing the traffic across the devices.

  - Preempt Enabled—Specifies whether the unit that is the preferred failover device for this failover group should become the active unit after rebooting.
  - Preempt Delay—Specifies the number of seconds that the preferred failover device should wait after rebooting before taking over as the active unit for this failover group. The range is between 0 and 1200 seconds.
  - Interface Policy—Specifies either the number of monitored interface failures or the percentage of failures that are allowed before the group fails over. The range is between 1 and 250 failures or 1 and 100 percent.
  - Interface Poll Time—Specifies the amount of time between polls among interfaces. The range is between 1 and 15 seconds.
  - Replicate HTTP—Identifies whether Stateful Failover should copy active HTTP sessions to the standby firewall for this failover group. If you do not allow HTTP replication, then HTTP connections are disconnected at failover. Disabling HTTP replication reduces the amount of traffic on the state link. This setting overrides the HTTP replication setting on the Setup tab.

- Add—Displays the Add Failover Group dialog box. This button is only enabled if less than 2 failover groups exist. See Add/Edit Failover Group for more information.
- Edit—Displays the Edit Failover Group dialog box for the selected failover group. See Add/Edit Failover Group for more information.
- Delete—Removes the currently selected failover group from the failover groups table. This button is only enabled if the last failover group in the list is selected.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

For More Information

For more information about failover in general, see Understanding Failover.

Add/Edit Failover Group

System > Configuration > Failover > Active/Active > Add/Edit Failover Group

Use the Add/Edit Failover Group dialog box to define failover groups for an Active/Active failover configuration.

Fields

- Preferred Role—Specifies the unit in the failover pair, primary or secondary, on which the failover group appears in the active state. You can have both failover groups be in the active state on a single unit in the pair, with the other unit containing the failover groups in the standby state. However, a more typical configuration is to assign each failover group a different role preference to make each one active on a different unit, balancing the traffic across the devices.

- Preempt after booting with optional delay of—Checking this check box causes the unit that is the preferred failover device for a failover group to become the active unit after rebooting. Checking this check box also enables the Preempt after booting with optional delay of field in which you can specify a period of time that the device should wait before becoming the active unit.
- Preempt after booting with optional delay of—Specifies the number of seconds that a unit should wait after rebooting before taking over as the active unit for any failover groups for which it is the preferred failover device. The range is between 0 and 1200 seconds.
- Interface Policy—Contains the fields for defining the policy for failover when monitoring detects an interface failure. These settings override any interface policy settings on the Criteria tab.
  - Number of failed interfaces that triggers failover—When the number of failed monitored interfaces exceeds the value you set with this command, then the security appliance fails over. The range is between 1 and 250 failures.
  - Percentage of failed interfaces that triggers failover—When the number of failed monitored interfaces exceeds the percentage you set with this command, then the security appliance fails over.
- Poll time interval for monitored interfaces—The amount of time between polls among interfaces. The range is between 1 and 15 seconds.
- Enable HTTP replication—Checking this check box enables Stateful Failover to copy active HTTP sessions to the standby firewall. If you do not allow HTTP replication, then HTTP connections are disconnected at failover. Disabling HTTP replication reduces the amount of traffic on the state link. This setting overrides the HTTP replication setting on the Setup tab.
- MAC Addresses—Lists physical interfaces on the security appliance for which an active and standby virtual MAC address has been configured.
  - Physical Interface—Displays the physical interface for which failover virtual MAC addresses are configured.
  - Active MAC Address—Displays the MAC address for the interface and failover group on the unit where the failover group is active.
  - Standby MAC Address—Displays the MAC address for the interface and failover group on the unit where the failover group is in the standby state.
- Add—Displays the Add Interface MAC Address dialog box. You cannot assign virtual MAC addresses to the LAN failover and Stateful Failover interfaces. See Add/Edit Interface MAC Address for more information.
- Edit—Displays the Edit Interface MAC Address dialog box for the selected interface. See Add/Edit Interface MAC Address for more information.
- Delete—Removes the currently selected interface from the MAC addresses table. There is no confirmation or undo.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

For More Information

For more information about failover in general, see Understanding Failover.

Add/Edit Interface MAC Address

System > Configuration > Failover > Active/Active > Add/Edit Failover Group > Add/Edit MAC Address

Use the Add/Edit Interface MAC Address dialog box to define the active and standby virtual MAC addresses for the interfaces in a failover group. If you do not specify a virtual MAC address for an interface, the interface is given a default virtual MAC address as follows:

- Active unit default MAC address: 00a0.c9*physical_port_number.failover_group_id*01.
- Standby unit default MAC address: 00a0.c9:*physical_port_number.failover_group_id*02.

Note If you have more than one Active/Active failover pair on the same network, it is possible to have the same default virtual MAC addresses assigned to the interfaces on one pair as are assigned to the interfaces of the other pairs because of the way the default virtual MAC addresses are determined. To avoid having duplicate MAC addresses on your network, make sure you assign each physical interface a virtual active and standby MAC address.

These MAC addresses override the physical MAC addresses for the interface.

Fields

- Physical Interface—Specifies the physical interface for which you are defining failover virtual MAC addresses. Because the MAC addresses do not change for the LAN failover and Stateful Failover interfaces during failover, you cannot choose these interfaces.
- MAC Addresses—Contains the fields for specifying the active and standby virtual MAC addresses for the interface.
  - Active Interface—Specifies the MAC address for the interface and failover group on the unit where the failover group is active. Each interface may have up to two MAC addresses, one for each failover group, which override the physical MAC address. Enter the MAC address in hexadecimal format (for example, 0123.4567.89AB).
  - Standby Interface—Specifies the MAC address for the interface and failover group on the unit where the failover group is in the standby state. Each interface may have up to two MAC addresses, one for each failover group, which override the physical MAC address. Enter the MAC address in hexadecimal format (for example, 0123.4567.89AB).

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

For More Information

For more information about failover in general, see Understanding Failover.

Failover > MAC Addresses Tab

System > Configuration > Failover > MAC Addresses Tab

The MAC Addresses tab lets you configure the virtual MAC addresses for the interfaces in an Active/Standby failover pair.

In Active/Standby failover, the MAC addresses for the primary unit are always associated with the active IP addresses. If the secondary unit boots first and becomes active, it uses the burned-in MAC address for its interfaces. When the primary unit comes online, the secondary unit obtains the MAC addresses from the primary unit. The change can disrupt network traffic.

You can configure virtual MAC addresses for each interface to ensure that the secondary unit uses the correct MAC addresses when it is the active unit, even if it comes online before the primary unit. If you do not specify virtual MAC addresses, then the failover pair uses the burned-in NIC address as the MAC address.

Note You cannot configure a virtual MAC address for the failover or state links. The MAC and IP addresses for those links do not change during failover.

In Active/Active failover, the MAC addresses configured on this tab are not in effect. Instead, the MAC addresses defined in the failover groups are used.

Fields

- MAC Addresses—Lists physical interfaces on the security appliance for which an active and standby virtual MAC address has been configured.
  - Physical Interface—Identifies the physical interface for which failover virtual MAC addresses are configured.
  - Active MAC Address—Identifies the MAC address on the active security appliance (usually primary).
  - Standby MAC Address—Identifies the MAC address on the standby security appliance (usually secondary).
- Add—Displays the Add/Edit Interface MAC Address dialog box.
- Edit—Displays the Add/Edit Interface MAC Address dialog box for the selected interface.
- Delete—Removes the currently selected interface from the MAC addresses table. There is no confirmation or undo.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

For More Information

For more information about failover in general, see Understanding Failover.

Add/Edit Interface MAC Address

System > Configuration > Failover > MAC Addresses > Add/Edit Interface MAC Address

Use the Add/Edit Interface MAC Address dialog box to define the active and standby virtual MAC addresses for an interface.

Fields

- Physical Interface—Specifies the physical interface for which you are defining failover virtual MAC addresses. Because the MAC addresses do not change for the LAN failover and Stateful Failover interfaces during failover, you cannot choose these interfaces.
- MAC Addresses—Contains the fields for specifying the active and standby virtual MAC addresses for the interface.
  - Active Interface—Specifies the MAC address of the interface on the active security appliance (usually primary). Enter the MAC address in hexadecimal format (for example, 0123.4567.89AB).
  - Standby Interface—Specifies the MAC address of the interface on the standby security appliance (usually secondary). Enter the MAC address in hexadecimal format (for example, 0123.4567.89AB).

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

For More Information

For more information about failover in general, see Understanding Failover.

CHAPTER 13

Configuring Logging

The Logging feature lets you enable logging and specify how log information is handled. The Log viewing feature lets you view system log messages in real-time. For a description of the Log viewing feature, see Chapter 37, "Monitoring System Log Messages.".

About Logging

The security appliance supports the generation of an audit trail of system log messages that describe its activities (for example, what kinds of network traffic has been allowed and denied) and enables you to configure system logging.

All system log messages have a default severity level. You can reassign a message to a new severity level, if necessary. When you choose a severity level, logging messages from that level and lower levels are generated. Messages from a higher level are not included. The higher the severity level, the more messages are included. For more information about logging and system log messages, see *Cisco Security Appliance Logging Configuration and System Log Messages*.

Security Contexts in Logging

Each security context includes its own logging configuration and generates its own messages. If you log in to the system or admin context, and then change to another context, messages you view in your session are only those that are related to the current context.

System log messages that are generated in the system execution space, including failover messages, are viewed in the admin context along with messages generated in the admin context. You cannot configure logging or view any logging information in the system execution space.

You can configure the security appliance to include the context name with each message, which helps you differentiate context messages that are sent to a single syslog server. This feature also helps you to determine which messages are from the admin context and which are from the system; messages that originate in the system execution space use a device ID of **system**, and messages that originate in the admin context use the name of the admin context as the device ID. To use the device ID, see Advanced Syslog Configuration, page 13-6.

Using Logging

After you have enabled logging, you can do the following:

Step 1 In the Logging Setup pane, configure the logging parameters. For more information, see Logging Setup, page 13-2.

Step 2 In the Syslog Setup pane, set the facility code to be included in system log messages that are sent to syslog servers, specify that a timestamp is included in each message, view the severity levels for messages, modify the severity level for messages, and suppress messages. For more information, see Syslog Setup, page 13-4.

Step 3 In the E-Mail Setup pane, specify system log messages to be sent by e-mail for notification purposes. For more information, see Syslog Setup, page 13-4.

Step 4 In the Event Lists pane, create custom lists of events that specify which messages should be logged; these lists are then used when you set up log filters. For more information, see Event Lists, page 13-8.

Step 5 In the Logging Filters pane, specify the criteria that should be used to filter the messages sent to each log destination. The criteria you use for creating filters are severity level, message class, message ID, or events lists. For more information, see Logging Filters, page 13-11.

Step 6 In the Rate Limit pane, limit the number of messages that can be generated in a specified time interval. For more information, see Rate Limit, page 13-15.

Step 7 In the Syslog Server pane, specify one or more syslog servers to which the security appliance sends system log messages. For more information, see Syslog Servers, page 13-18.

Logging Setup

SupportedUserRoles

Configuration > Properties > Logging > Logging Setup

The Logging Setup pane lets you enable system logging on the security appliance and lets you specify general logging parameters, including whether standby units can take over logging, whether to send debug messages, and whether to use the EMBLEM format. It also lets you change default settings for the internal log buffer and the security appliance logging queue.

Fields

- Enable logging—Turns on logging for the main security appliance.
- Enable logging on the failover standby unit—Turns on logging for the standby security appliance, if available.
- Send debug messages as syslogs—Redirects all debug trace output to system logs. The system log message does not appear in the console if this option is enabled. Therefore, to view debug messages, you must have logging enabled at the console and have it configured as the destination for the debug system log message number and severity level. The system log message number used is 711001. The default severity level for this system log message is debug.
- Send syslogs in EMBLEM format—Enables EMBLEM format so that it is used for all log destinations except syslog servers.
- Buffer Size—Specifies the size of the internal log buffer to which system log messages are saved if the logging buffer is enabled. When the buffer fills up, it will be overwritten unless you choose to enable saving of the logs to an FTP server or to internal Flash memory. The default buffer size is 4096 bytes. The range is 4096 to 1048576.

- Save Buffer To FTP Server—To save the buffer contents to the FTP server before it is overwritten, check this . To remove the FTP configuration, uncheck this box.
- Configure FTP Settings—Identifies the FTP server and configures the FTP parameters used to save the buffer content.
- Save Buffer To Flash—To save the buffer contents to internal Flash memory before it is overwritten, check this .

Note This option is only available in routed or transparent single mode.

- Configure Flash Usage—Specifies the maximum space to be used in internal Flash memory for logging and the minimum free space to be preserved (in KB). Enabling this option creates a directory called “syslog” on the device disk in which messages are stored.

Note This option is only available in routed or transparent single mode.

- security appliance Logging Queue Size—Specifies the queue size for system logs that are to be viewed in security appliance.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

- See Configure FTP Settings, page 13-3.
- See Configure Logging Flash Usage, page 13-4.

Configure FTP Settings

Configuration > Features > Properties > Logging > Logging Setup > Configure FTP Settings

The Configure FTP Settings dialog box lets you specify the configuration for the FTP server that is used to save the buffer contents.

Fields

- Enable FTP client—Enables the configuration of the FTP client.
- Server IP Address—IP address of the FTP server.
- Path—Directory path on the FTP server to store the saved file.
- Username—Username to log in to the FTP server.
- Password—Password associated with the username to log in to the FTP server.
- Confirm Password—Confirms the password.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Configure Logging Flash Usage

Configuration > Properties > Logging > Logging Setup > Configure Logging Flash Usage

The Configure Logging Flash Usage dialog box lets you specify the limits for saving buffer contents to internal Flash memory.

Fields

- Maximum Flash to Be Used by Logging—Specifies the maximum amount of internal Flash memory that can be used for logging (in KB).
- Minimum Free Space to Be Preserved—Specifies the amount of internal Flash memory that is preserved (in KB). When the internal Flash memory approaches that limit, new logs are not saved.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

Syslog Setup

Configuration > Properties > Logging > Syslog Setup

The Syslog Setup pane lets you set the facility code to include in messages destined for syslog servers and determine whether system log messages should include the timestamp. It also lets you change message severity levels and suppress messages you do not want to be logged.

Fields

- Facility code to include in syslogs—Specifies a system log facility for syslog servers to use as a basis to file messages. The default is LOCAL(4)20, which is what most UNIX systems expect. However, because your network devices share the eight available facilities, you might need to change this value for system logs.
- Include timestamp in syslogs—Includes date and time in every system log message sent.
- Syslog ID Setup—Selects the information to be displayed in the Syslog ID Table. Options are defined as follows:

  - Show all syslog IDs—Specifies that the syslog ID table should display the entire list of system log message IDs.
  - Show suppressed syslog IDs—Specifies that the syslog ID table should display only those system log message IDs that have been explicitly suppressed.
  - Show syslog IDs with changed logging—Specifies that the syslog ID table should display only those system log message IDs with severity levels that have changed from their default values.
  - Show syslog IDs that are suppressed or with a changed logging level—Specifies that the syslog ID table should display only those system log message IDs with severity levels that have been modified and the IDs of system log messages that have been explicitly suppressed.
- Syslog ID Table—*Display only.* Shows the list of system log messages based on the setting in the Syslog ID Table View. Select individual messages or ranges of message IDs that you want to modify. You can either suppress the selected message IDs or modify their severity levels. To select more than one message ID in the list, click the first ID in the range and Shift-click the last ID in the range.
- Advanced—Lets you configure system log messages to include a device ID.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

- See Edit Syslog ID Settings, page 13-5.
- See Advanced Syslog Configuration, page 13-6.

Edit Syslog ID Settings

Configuration > Properties > Logging > Syslog Setup > Edit Syslog ID Settings

The Edit Syslog ID Settings dialog box lets you modify the severity level of the selected system log messages or specify that the selected system log messages should be suppressed.

Fields

- Syslog ID(s)—This text area is read-only. The values displayed in this area are determined by the entries selected in the Syslog ID Table located in the Syslog Setup pane.
- Suppress Message(s)—Check this to suppress messages for the system log message ID(s) displayed in the Syslog ID(s) list.
- Logging Level—Choose the severity level of messages to be sent for the system log message ID(s) displayed in the Syslog ID(s) list. Levels are defined as follows:
  - Emergency (level 0, system unusable)
  - Alert (level 1, immediate action needed)
  - Critical (level 2, critical condition)

- Error (level 3, error condition)
- Warning (level 4, warning condition)
- Notification (level 5, normal but significant condition)
- Informational (level 6, informational message only)
- Debugging (level 7, appears during debugging only)

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Advanced Syslog Configuration

Configuration > Properties > Logging > Syslog Setup > Advanced > Advanced Syslog Configuration

You can configure the security appliance to include a device ID in non-EMBLEM-format system log messages. You can specify only one type of device ID for the system log messages. The device ID can be the hostname of the FWSM, an interface IP address, the context, or a text string.

The Advanced Syslog Configuration dialog box lets you determine whether system log messages should include a device ID. If this feature is enabled, the device ID is included in all non-EMBLEM formatted system log messages.

Fields

- Enable Syslog Device ID—Specifies that a device ID should be included in all non-EMBLEM formatted system log messages.
- Hostname—Specifies that the hostname is used as the device ID.
- IP Address—Specifies the IP address of the interface that is used as the device ID.
  - Interface Name—Specifies the interface name corresponding to the specified IP address.
- String—Specifies that a user-defined string is used as the device ID.
  - User-defined ID—Specifies an alphanumeric user-defined string.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

E-Mail Setup

Configuration > Properties > Logging > E-Mail Setup

The E-Mail Setup pane lets you set up a source e-mail address as well as a list of recipients for specified system log messages to be sent as e-mail messages for notification purposes. You can filter the system log messages sent to a destination e-mail address by severity level. The table shows which entries have been set up.

The system log message severity level used to filter messages for a destination e-mail address is the higher of the severity level selected in this section compared to the global filter set for all e-mail recipients in the Logging Filters pane.

The system log message severity filter used for the destination e-mail address causes messages of the specified severity level and higher to be sent. The global filter specified in the Logging Filters pane is also applied to each e-mail recipient.

Fields

- Source E-Mail address—Specifies the e-mail address that is used as the source address for system log messages sent as e-mail messages.
- Destination E-Mail Address—Specifies the e-mail address of the recipient of the specified system log messages.
- Syslog Severity—Specifies the severity level of the system log messages that should be sent to this recipient. Messages with the specified severity level and higher are sent.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

- See Add/Edit E-Mail Recipients, page 13-7.
- See Logging Filters, page 13-11.

Add/Edit E-Mail Recipients

Configuration > Properties > Logging > E-Mail Setup > Add/Edit E-Mail Recipient

The Add/Edit E-Mail Recipient dialog box lets you set up a destination e-mail address for a particular severity of system log messages to be sent as e-mail messages.

The severity level used to filter messages for the destination e-mail address is the higher of the severity level selected in this section compared to the global filter set for all e-mail recipients in the Logging Filters pane.

Fields

- Destination E-Mail Address—Specifies the e-mail address of the recipient of selected system log messages.
- Syslog Severity—Specifies the severity level of the system log messages sent to this recipient.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Event Lists

Configuration > Properties > Logging > Event Lists

The Event Lists pane lets you create custom lists of events that are used to select which system log messages are sent to a particular destination. After you enable logging and configure the logging parameters using the Logging Setup pane, create one or more lists of events on the Event Lists pane. Use these lists on the Logging Filters pane to specify a logging destination for each list of events.

You can use three criteria to define an event list:

- Message Class
- Severity
- Message ID.

A message class is a group of system log messages related to a security appliance feature that enables you to specify an entire class of messages rather than specifying each message individually. For example, use the auth class to select all system log messages that are related to user authentication.

Severity classifies system log messages based on the relative importance of the event in the normal functioning of the network. The highest severity is emergency, which means the resource is no longer available. The lowest severity is debugging, which provides detailed information about every network event.

The message ID is a numeric value that uniquely identifies each message. You can use the message ID in an event list to identify a range of system log messages, such as 101001-101010.

Fields

- Name—Lists the name of the event list.
- Event Class/Severity—Lists the event class and the level of logging messages. Event classes include:
  - All—All event classes
  - auth—User Authentication
  - bridge—Transparent firewall
  - ca—PKI Certification Authority

  - config—Command Interface
  - ha—Failover
  - ids—Intrusion Detection System
  - ip—IP Stack
  - np—Network Processor
  - ospf—OSPF Routing
  - rip—RIP Routing
  - rm—Resource Manager
  - session—User Session
  - snmp—SNMP
  - sys—System

  Severity levels include the following:

  - Emergency (level 0, system unusable)
  - Alert (level 1, immediate action needed)
  - Critical (level 2, critical condition)
  - Error (level 3, error condition)
  - Warning (level 4, warning condition)
  - Notification (level 5, normal but significant condition)
  - Informational (level 6, informational message only)
  - Debugging (level 7, appears during debugging only)
- Message IDs—Lists a system log message ID or range of IDs (for example, 101001-101010) to include in the filter.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

- See Add/Edit Event List, page 13-9.
- See Add/Edit Syslog Message ID Filter, page 13-11.
- See Logging Filters, page 13-11.

Add/Edit Event List

Configuration > Properties > Logging > Event Lists > Add/Edit Event List (You can get to this dialog box through various paths.)

The Add/Edit Event List dialog box lets you create or edit an event list that you can use to specify which messages should be sent to a log destination. You can create event lists that filter messages according to message class and severity, or by message ID.

A message class is a group of system log messages related to a security appliance feature. When creating an event list, you can specify an entire class of messages rather than specifying each message individually. For example, use the auth class to select all system log messages that are related to user authentication.

Severity defines system log messages based on the relative importance of the event in the normal functioning of the network. The highest severity is emergency, which means the resource is no longer available. The lowest severity is debugging, which provides detailed information about every network event.

The message ID is a numeric value that uniquely identifies each message. You can use the message ID in an event list to identify a range of system log messages, such as 101001-101010.

Fields

- Name—Enter the name of the event list.
- Event Class—Lists the event class. Event classes include:
  - All—All event classes
  - auth—User Authentication
  - bridge—Transparent firewall
  - ca—PKI Certification Authority
  - config—Command Interface
  - ha—Failover
  - ips—Intrusion Protection Service
  - ip—IP Stack
  - np—Network Processor
  - ospf—OSPF Routing
  - rip—RIP Routing
  - rm—Resource Manager
  - session—User Session
  - snmp—SNMP
  - sys—System
- Severity—Lists the level of logging messages. Severity levels include the following:
  - Emergency (level 0, system unusable)
  - Alert (level 1, immediate action needed)
  - Critical (level 2, critical condition)
  - Error (level 3, error condition)
  - Warning (level 4, warning condition)
  - Notification (level 5, normal but significant condition)
  - Informational (level 6, informational message only)
  - Debugging (level 7, appears during debugging only)

- Message IDs Filters—Lists a system log message ID or range of system log message IDs, such as 101001-101010, to include in the filter.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Syslog Message ID Filter

Configuration > Properties > Logging > Event Lists > Add/Edit Event List > Add/Edit Syslog Message ID Filter

The Add/Edit Syslog Message ID Filter dialog box lets you specify one or more system log message IDs to be included in the event list.

Fields

- Message IDs—Specify a system log message ID or range of IDs to be logged. Use a hyphen to specify a range (for example, 101001-101010).

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Logging Filters

Configuration > Properties > Logging > Logging Filters

The Logging Filters pane lets you apply message filters to a log destination. Filters applied to a log destination select the messages that are sent to that destination.

You can filter messages according to message class and severity level, or use an event list that you can create on the Event Lists pane.

Fields

- Logging Destination—Lists the name of the logging destination to which you can apply a filter. Logging destinations are as follows:
  - Console
  - Security appliance

  - Syslog Servers
  - SNMP Trap
  - E-Mail
  - Internal Buffer
  - Telnet Sessions
- Syslogs From All Event Classes—Lists the severity or the event list to use to filter messages for the log destination, or whether logging is disabled for all event classes.
- Syslogs From Specific Event Classes—Lists the event class to use to filter messages for that log destination.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

- See Edit Logging Filters, page 13-12.
- See Add/Edit Syslog Message ID Filter, page 13-11.
- See Add/Edit Class and Severity Filter, page 13-13.
- See Event Lists, page 13-8.

Edit Logging Filters

Configuration > Properties > Logging > Logging Filters > Edit Logging Filters

The Edit Logging Filters dialog box lets you apply filters to each log destination, edit filters already applied to a log destination, or disable filters for the log destination.

You can filter messages according to message class and severity level, or use an event list that you can create on the Event Lists pane.

Fields

- Logging Destination—Specifies the logging destination for this filter.
- Filter on severity—Filters system log messages according to their severity level.
  - Filter on severity—Specifies the level of system log messages on which to filter.
- Use event list—Specifies that an event list will be used for this filter.
  - Use event—Specifies the event list to use.
- New—Lets you add a new event list.
- Disable logging from all event classes—Disables all logging to the selected destination.
- Event Class—Specifies the event class. Event classes include:

  - All—All event classes
  - auth—User Authentication
  - bridge—Transparent firewall
  - ca—PKI Certification Authority
  - config—Command Interface
  - ha—Failover
  - ids—Intrusion Detection System
  - ip—IP Stack
  - np—Network Processor
  - ospf—OSPF Routing
  - rip—RIP Routing
  - rm—Resource Manager
  - session—User Session
  - snmp—SNMP
  - sys—System
- Severity—Specifies the level of logging messages. Severity levels include:
  - Emergency (level 0, system unusable)
  - Alert (level 1, immediate action needed)
  - Critical (level 2, critical condition)
  - Error (level 3, error condition)
  - Warning (level 4, warning condition)
  - Notification (level 5, normal but significant condition)
  - Informational (level 6, informational message only)
  - Debugging (level 7, appears during debugging only)

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Class and Severity Filter

...Add Event List > Add/Edit Class and Severity Filter (You can get to this dialog box through various paths.)

The Add/Edit Class and Severity Filter dialog box lets you specify a message class and severity level to be used to filter messages.

A message class is a group of system log messages related to a security appliance feature. When creating an event list, you can specify an entire class of messages rather than specifying each message individually. For example, use the auth class to select all of the system log messages that are related to user authentication.

Severity defines system logs based on the relative importance of the event in the normal functioning of the network. The highest severity is emergency, which means the resource is no longer available. The lowest severity is debugging, which provides detailed information about every network event.

Fields

- Event Class—Specifies the event class. Event classes include:
  - All—All event classes
  - auth—User Authentication
  - bridge—Transparent firewall
  - ca—PKI Certification Authority
  - config—Command Interface
  - ha—Failover
  - ids—Intrusion Detection System
  - ip—IP Stack
  - np—Network Processor
  - ospf—OSPF Routing
  - rip—RIP Routing
  - rm—Resource Manager
  - session—User Session
  - snmp—SNMP
  - sys—System
- Severity—Specifies the level of logging messages. Severity levels include:
  - Emergency (level 0, system unusable)
  - Alert (level 1, immediate action needed)
  - Critical (level 2, critical condition)
  - Error (level 3, error condition)
  - Warning (level 4, warning condition)
  - Notification (level 5, normal but significant condition)
  - Informational (level 6, informational message only)
  - Debugging (level 7, appears during debugging only)

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Syslog Message ID Filter

Configuration > Properties > Logging > Logging Filters > Edit Logging Filters > Add Event List > Add/Edit Syslog Message ID Filter

The Add/Edit Syslog Message ID Filter dialog box lets you specify individual system log message IDs or ranges of IDs to include in the event list filter.

Fields

- Message IDs—Specifies the system log message ID or range of IDs. Use a hyphen to specify a range (for example, 101001-101010).

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Rate Limit

Configuration > Features > Properties > Logging > Rate Limit

The Rate Limit pane lets you specify the number of system log messages that the firewall can send. You must also enable logging using the Logging Setup pane. You can specify a rate limit for message logging levels or be more specific and limit the rate of a specific message. The rate level is applied to the severity level or to the message ID, not to a destination. Therefore, rate limits affect the volume of messages being sent to all configured destinations.

Fields

Rate limits for syslog logging levels

- Logging Level—Lists the message severity level. Levels are defined as follows:
  - Disabled (no logging)
  - Emergency (level 0, system unusable)
  - Alert (level 1, immediate action needed)
  - Critical (level 2, critical condition)
  - Error (level 3, error condition)

  - Warning (level 4, warning condition)
  - Notification (level 5, normal but significant condition)
  - Informational (level 6, informational message only)
  - Debugging (level 7, appears during debugging only)
- No of Messages—Displays the number of messages sent. To allow an unlimited number of messages, leave both the Number of Messages and Time Interval fields blank.
- Interval (Seconds)—Displays the interval, in seconds, used to limit how many messages at this logging level can be sent. To allow an unlimited number of messages, leave both the Number of Messages and Time Interval blank.
- Edit—Select a logging level from the table and click this button to open the Edit Rate Limit dialog box, where you can edit the properties of the selected logging level.
- **Individually rate-limited syslog messages**
  - Syslog ID—Displays the ID for the system log message that is limited.
  - Logging Level—Displays the message severity level. For a list of severity levels, see Rate limits for syslog logging levels, page 13-15.
  - No of Messages—Displays the maximum number of messages that can be sent in the specified time interval.
  - Interval (Seconds)—Displays the interval, in seconds, used to limit the system log message.
  - Add—Click this button to limit the rate of a specific message.
- Apply—Sends changes to the firewall and applies them to the running configuration. Use the File menu to write a copy of the running configuration to internal Flash memory, a TFTP server, or a failover standby firewall unit.
- Reset—Discards changes and reverts values to those displayed when it was opened or the last time Refresh was clicked while open.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

- See Edit Rate Limit for Syslog Logging Level, page 13-16.
- See Add/Edit Rate Limit for Syslog Message, page 13-17.

Edit Rate Limit for Syslog Logging Level

Configuration > Features > Properties > Logging > Edit Rate Limit for Syslog Logging Level

The Edit Rate Limit for Syslog Logging Level box lets you limit the number of messages the firewall can send in a specified time interval.

Fields

Rate limit for syslog logging levels

- Logging Level—Displays the selected message severity level. If you are modifying a specific message ID rate limit, you may specify the logging level. Levels are defined as follows:
  - Disabled (no logging)
  - Emergency (level 0, system unusable)
  - Alert (level 1, immediate action needed)
  - Critical (level 2, critical condition)
  - Error (level 3, error condition)
  - Warning (level 4, warning condition)
  - Notification (level 5, normal but significant condition)
  - Informational (level 6, informational message only)
  - Debugging (level 7, appears during debugging only)
- No of Messages—Specifies the maximum number of messages at this logging level that can be sent.
- Time Interval (seconds)—Specifies the amount of time, in seconds, used to limit the messages at this logging level.
- OK—Accepts changes and returns to the previous pane.
- Cancel—Discards changes and returns to the previous pane.
- Help—Provides more information.
- Reset—Discards changes and reverts values to those displayed when it was opened or the last time Refresh was clicked while open.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Rate Limit for Syslog Message

Configuration > Features > Properties > Logging > Rate Limit > Add/Edit Rate Limit for Syslog Message

The Add/Edit Rate Limit for Syslog Message dialog box lets you assign rate limits to a specific system log message.

Fields

- Syslog Message ID—Specifies the message ID of the system log message you want to limit.
- Number of Messages—Specifies the maximum number of times this message can be sent in the specified time interval.

- Time Interval—Specifies the amount of time, in seconds, used to limit the specified message.

Note To allow an unlimited number of messages, leave both Number of Messages and Time Interval blank.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Syslog Servers

Configuration > Properties > Logging > Syslog Servers

The Syslog Servers pane lets you specify the syslog servers to which the security appliance should send system log messages. To use the syslog server(s) you define, you must enable logging using the Logging Setup pane and set up the appropriate destinations in the Logging Filters pane.

Note You can set up a maximum of four syslog servers per context.

Fields

- Interface—Displays the interface used to communicate with the syslog server.
- IP Address—Displays the IP address of the interface that will be used to communicate with the syslog server.
- Protocol/Port—Displays the protocol and port that the syslog server will use to communicate with the security appliance.
- EMBLEM—Specifies whether to log messages in Cisco EMBLEM format (available only if UDP is selected in the Protocol/Port).
- Queue Size—Specifies the number of messages that are allowed to be queued on the security appliance if any syslog server is busy. A zero value means an unlimited number of messages may be queued.
- Allow user traffic to pass when TCP syslog server is down—Specifies whether to restrict all traffic if any syslog server is down.
- Deny connection upon queue full—Specifies whether to allow connections when the queue fills (that is, when it reaches the limit set in the Queue Size).

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

- See Add/Edit Syslog Server, page 13-19.
- See Logging Setup, page 13-2.
- See Logging Filters, page 13-11.

Add/Edit Syslog Server

Configuration > Properties > Logging > Syslog Servers > Add/Edit Syslog Server

The Add/Edit Syslog Server dialog box lets you add or edit the syslog servers to which the security appliance sends system log messages. To use the syslog server(s) you define, you must enable logging in the Logging Setup pane and set up the appropriate filters for log destinations in the Logging Filters pane.

Note You can set up a maximum of four syslog servers per context.

Fields

- Interface—Specifies the interface used to communicate with the syslog server.
- IP Address—Specifies the IP address used to communicate with the syslog server.
- Protocol—Displays the protocol (either TCP or UDP) used by the syslog server to communicate with the security appliance.
- Port—Specifies the port used by the syslog server to communicate with the security appliance.
- Log messages in Cisco EMBLEM format (UDP only)—Specifies whether to log messages in Cisco EMBLEM format (available only if UDP is selected in the Protocol).

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Configuring Dynamic And Static Routing

The Routing area lets you edit a static route to ensure that the security appliance correctly forwards network packets destined to the host or network. You can also use a static route to override any dynamic routes that are discovered for this host or network by specifying a static route with a lower metric than the discovered dynamic routes. To create a static route for a host or network, you must define the IP address and metric for the hop gateway to which the security appliance will forward packets destined to the selected host or network. You can also define multiple static routes for a host or network.

This section contains the following topics:

Dynamic Routing

The Dynamic Routing area contains the following topics:

OSPF

OSPF is an interior gateway routing protocol that uses link states rather than distance vectors for path selection. OSPF propagates link-state advertisements rather than routing table updates. Because only LSAs are exchanged, rather than entire routing tables, OSPF networks converge more quickly than RIP networks.

OSPF supports MD5 and clear text neighbor authentication. Authentication should be used with all routing protocols when possible because route redistribution between OSPF and other protocols (like RIP) can potentially be used by attackers to subvert routing information.

If NAT is used, if OSPF is operating on public and private areas, and if address filtering is required, then you need to run two OSPF processes—one process for the public areas and one for the private areas.

A router that has interfaces in multiple areas is called an Area Border Router (ABR). A router that acts as a gateway to redistribute traffic between routers using OSPF and routers using other routing protocols is called an Autonomous System Boundary Router (ASBR).

An ABR uses LSAs to send information about available routes to other OSPF routers. Using ABR type 3 LSA filtering, you can have separate private and public areas with the security appliance acting as an ABR. Type 3 LSAs (inter-area routes) can be filtered from one area to other. This lets you use NAT and OSPF together without advertising private networks.

Note Only type 3 LSAs can be filtered. If you configure the security appliance as an ASBR in a private network, it will send type 5 LSAs describing private networks, which will get flooded to the entire AS including public areas.

If NAT is employed but OSPF is only running in public areas, then routes to public networks can be redistributed inside the private network, either as default or type 5 AS External LSAs. However, you need to configure static routes for the private networks protected by the security appliance. Also, you should not mix public and private networks on the same security appliance interface.

You can have two OSPF routing processes and one RIP routing process running on the security appliance at the same time.

For more information about enabling and configuring OSPF, see the following:

- Setup
- Interface
- Static Neighbor
- Virtual Link
- Filtering
- Redistribution
- Summary Address

Setup

Configuration > Routing > Dynamic Routing > OSPF > Setup

The Setup pane lets you enable OSPF processes, configure OSPF areas and networks, and define OSPF route summarization.

For more information about configuring these areas, see the following:

- Setup > Process Instances Tab
- Setup > Area/Networks Tab
- Setup > Route Summarization Tab

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Setup > Process Instances Tab

Configuration > Routing > Dynamic Routing > OSPF > Setup > Process Instances Tab

You can enable up to two OSPF process instances. Each OSPF process has its own associated areas and networks.

Fields

- OSPF Process 1 and 2 areas—Each area contains the settings for a specific OSPF process.
- Enable this OSPF Process—Check the check box to enable an OSPF process. You cannot enable an OSPF process if you have RIP enabled on the security appliance. Uncheck this check box to remove the OSPF process.
- OSPF Process ID—Enter a unique numeric identifier for the OSPF process. This process ID is used internal and does not need to match the OSPF process ID on any other OSPF devices. Valid values are from 1 to 65535.
- Advanced—Opens the Edit OSPF Process Advanced Properties dialog box, where you can configure the Router ID, Adjacency Changes, Administrative Route Distances, Timers, and Default Information Originate settings.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Edit OSPF Process Advanced Properties

Configuration > Routing > Dynamic Routing > OSPF > Setup > Process Instances > Edit OSPF Process Advanced Properties

You can edit process-specific settings, such as the Router ID, Adjacency Changes, Administrative Route Distances, Timers, and Default Information Originate settings, in the Edit OSPF Process Advanced Properties dialog box.

Fields

- OSPF Process—Displays the OSPF process you are configuring. You cannot change this value.
- Router ID—To used a fixed router ID, enter a router ID in IP address format in the Router ID field. If you leave this value blank, the highest-level IP address on the security appliance is used as the router ID.
- Ignore LSA MOSPF—Check this check box to suppress the sending of system log messages when the security appliance receives Type 6 (MOSPF) LSA packets. This setting is unchecked by default.
- RFC 1583 Compatible—Check this check box to calculate summary route costs per RFC 1583. Uncheck this check box to calculate summary route costs per RFC 2328. To minimize the chance of routing loops, all OSPF devices in an OSPF routing domain should have RFC compatibility set identically.This setting is selected by default.

- Adjacency Changes—Contains settings that define the adjacency changes that cause system log messages to be sent.
  - Log Adjacency Changes—Check this check box to cause the security appliance to send a system log message whenever an OSPF neighbor goes up or down. This setting is selected by default.
  - Log Adjacency Changes Detail—Check this check box to cause the security appliance to send a system log message whenever any state change occurs, not just when a neighbor goes up or down. This setting is unchecked by default.
- Administrative Route Distances—Contains the settings for the administrative distances of routes based on the route type.
  - Inter Area—Sets the administrative distance for all routes from one area to another. Valid values range from 1 to 255. The default value is 100.
  - Intra Area—Sets the administrative distance for all routes within an area. Valid values range from 1 to 255. The default value is 100.
  - External—Sets the administrative distance for all routes from other routing domains that are learned through redistribution. Valid values range from 1 to 255. The default value is 100.
- Timers—Contains the settings used to configure LSA pacing and SPF calculation timers.
  - SPF Delay Time—Specifies the time between when OSPF receives a topology change and when the SPF calculation starts. Valid values range from 0 to 65535. The default value is 5.
  - SPF Hold Time—Specifies the hold time between consecutive SPF calculations.Valid values range from 1 to 65534. The default value is 10.
  - LSA Group Pacing—Specifies the interval at which LSAs are collected into a group and refreshed, checksummed, or aged. Valid values range from 10 to 1800. The default value is 240.
- Default Information Originate—Contains the settings used by an ASBR to generate a default external route into an OSPF routing domain.
  - Enable Default Information Originate—Check this check box to enable the generation of the default route into the OSPF routing domain.
  - Always advertise the default route—Check this check box to always advertise the default route. This option is unchecked by default.
  - Metric Value—Specifies the OSPF default metric. Valid values range from 0 to 16777214. The default value is 1.
  - Metric Type—Specifies the external link type associated with the default route advertised into the OSPF routing domain. Valid values are 1 or 2, indicating a Type 1 or a Type 2 external route. The default value is 2.
  - Route Map—(Optional) The name of the route map to apply. The routing process generates the default route if the route map is satisfied.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Setup > Area/Networks Tab

Configuration > Routing > Dynamic Routing > OSPF > Setup > Area/Networks Tab

The Area/Networks tab displays the areas, and the networks they contain, for each OSPF process on the security appliance.

Fields

- Area/Networks—Displays information about the areas and the area networks configured for each OSPF process. Double-clicking a row in the table opens the Add/Edit OSPF Area dialog box for the selected area.
  - OSPF Process—Displays the OSPF process the area applies to.
  - Area ID—Displays the area ID.
  - Area Type—Displays the area type. The area type is one of the following values: Normal, Stub, NSSA.
  - Networks—Displays the area networks.
  - Authentication—Displays the type of authentication set for the area. The authentication type is one of the following values: None, Password, MD5.
  - Options—Displays any options set for the area type.
  - Cost—Displays the default cost for the area.
- Add—Opens the Add/Edit OSPF Area dialog box. Use this button to add a new area configuration.
- Edit—Opens the Add/Edit OSPF Area dialog box. Use this button to change the parameters of the selected area.
- Delete—Removes the selected area from the configuration.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit OSPF Area

Configuration > Routing > Dynamic Routing > OSPF > Setup > Area/Networks > Add/Edit OSPF Area

You define area parameters, the networks contained by the area, and the OSPF process associated with the area in the Add/Edit OSPF Area dialog box.

Fields

- OSPF Process—When adding a new area, choose the OSPF process ID for the OSPF process for which the area is being. If there is only one OSPF process enabled on the security appliance, then that process is selected by default. When editing an existing area, you cannot change the OSPF process ID.

- Area ID—When adding a new area, enter the area ID. You can specify the area ID as either a decimal number or an IP address. Valid decimal values range from 0 to 4294967295. You cannot change the area ID when editing an existing area.
- Area Type—Contains the settings for the type of area being configured.
  - Normal—Choose this option to make the area a standard OSPF area. This option is selected by default when you first create an area.
  - Stub—Choosing this option makes the area a stub area. Stub areas do not have any routers or areas beyond it. Stub areas prevent AS External LSAs (Type 5 LSAs) from being flooded into the stub area. When you create a stub area, you have the option of preventing summary LSAs (Type 3 and 4) from being flooded into the area by unchecking the Summary check box.
  - Summary—When the area being defined is a stub area, unchecking this check box prevents LSAs from being sent into the stub area. This check box is selected by default for stub areas.
  - NSSA—Choose this option to make the area a not-so-stubby area. NSSAs accept Type 7 LSAs. When you create a NSSA, you have the option of preventing summary LSAs from being flooded into the area by unchecking the Summary check box. You can also disable route redistribution by unchecking the Redistribute check box and enabling Default Information Originate.
  - Redistribute—Uncheck this check box to prevent routes from being imported into the NSSA. This check box is selected by default.
  - Summary—When the area being defined is a NSSA, unchecking this check box prevents LSAs from being sent into the stub area. This check box is selected by default for NSSAs.
  - Default Information Originate—Check this check box to generate a Type 7 default into the NSSA. This check box is unchecked by default.
  - Metric Value—Specifies the OSPF metric value for the default route. Valid values range from 0 to 16777214. The default value is 1.
  - Metric Type—The OSPF metric type for the default route. The choices are 1 (Type 1) or 2 (Type 2). The default value is 2.
- Area Networks—Contains the settings for defining an OSPF area.
  - Enter IP Address and Mask—Contains the settings used to define the networks in the area.

    IP Address—Enter the IP address of the network or host to be added to the area. Use 0.0.0.0 with a netmask of 0.0.0.0 to create the default area. You can only use 0.0.0.0 in one area.

    Netmask—Choose the network mask for the IP address or host to be added to the area. If adding a host, choose the 255.255.255.255 mask.
  - Add—Adds the network defined in the Enter IP Address and Mask area to the area. The added network appears in the Area Networks table.
  - Delete—Deletes the selected network from the Area Networks table.
  - Area Networks—Displays the networks defined for the area.

    IP Address—Displays the IP address of the network.

    Netmask—Displays the network mask for the network.
- Authentication—Contains the settings for OSPF area authentication.
  - None—Choose this option to disable OSPF area authentication. This is the default setting.
  - Password—Choose this option to use a clear text password for area authentication. This option is not recommended where security is a concern.
  - MD5—Choose this option to use MD5 authentication.

- Default Cost—Specify a default cost for the area. Valid values range from 0 to 65535. The default value is 1.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Setup > Route Summarization Tab

Configuration > Routing > Dynamic Routing > OSPF > Setup > Route Summarization Tab

In OSPF, an ABR will advertise networks in one area into another area. If the network numbers in an area are assigned in a way such that they are contiguous, you can configure the ABR to advertise a summary route that covers all the individual networks within the area that fall into the specified range. To define summary address for external routes being redistributed into an OSPF area, see Summary Address.

Fields

- Route Summarization—Displays information about route summaries defined on the security appliance. Double-clicking a row in the table opens the Add/Edit Route Summarization dialog box for the selected route summary.
  - OSPF Process—Displays the OSPF process ID for the OSPF process associated with the route summary.
  - Area ID—Displays the area associated with the route summary.
  - IP Address—Displays the summary address.
  - Network Mask—Displays the summary mask.
  - Advertise—Displays “yes” when the route summaries are advertised when they match the address/mask pair or “no” when route summaries are suppressed when they match the address/mask pair.
- Add—Opens the Add/Edit Route Summarization dialog box. Use this button to define a new route summarization.
- Edit—Opens the Add/Edit Route Summarization dialog box. Use this button to change the parameters of the selected route summarization.
- Delete—Removes the selected route summarization from the configuration.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Route Summarization

Configuration > Routing > Dynamic Routing > OSPF > Setup > Route Summarization > Add/Edit Route Summarization

Use the Add Route Summarization dialog box to add a new entry to the Route Summarization table. Use the Edit Route Summarization dialog box to change an existing entry.

Fields

- OSPF Process—Choose the OSPF process the route summary applies to. You cannot change this value when editing an existing route summary entry.
- Area ID—Choose the area ID the route summary applies to. You cannot change this value when editing an existing route summary entry.
- IP Address—Enter the network address for the routes being summarized.
- Network Mask—Choose one of the common network masks from the list or type the mask in the field.
- Advertise—Check this check box to set the address range status to “advertise”. This causes Type 3 summary LSAs to be generated. Uncheck this check box to suppress the Type 3 summary LSA for the specified networks. This check box is checked by default.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Filtering

Configuration > Routing > Dynamic Routing > OSPF > Filtering

The Filtering pane displays the ABR Type 3 LSA filters that have been configured for each OSPF process.

ABR Type 3 LSA filters allow only specified prefixes to be sent from one area to another area and restricts all other prefixes. This type of area filtering can be applied out of a specific OSPF area, into a specific OSPF area, or into and out of the same OSPF areas at the same time.

Benefits

OSPF ABR Type 3 LSA filtering improves your control of route distribution between OSPF areas.

Restrictions

Only Type-3 LSAs that originate from an ABR are filtered.

Fields

The Filtering table displays the following information. Double-clicking a table entry opens the Add/Edit Filtering Entry dialog box for the selected entry.

- OSPF Process—Displays the OSPF process associated with the filter entry.
- Area ID—Displays the ID of the area associated with the filter entry.
- Filtered Network—Displays the network address being filtered.
- Traffic Direction—Displays “Inbound” if the filter entry applies to LSAs coming in to an OSPF area or Outbound if it applies to LSAs coming out of an OSPF area.
- Sequence #—Displays the sequence number for the filter entry. When multiple filters apply to an LSA, the filter with the lowest sequence number is used.
- Action—Displays “Permit” if LSAs matching the filter are allowed or “Deny” if LSAs matching the filter are denied.
- Lower Range—Displays the minimum prefix length to be matched.
- Upper Range—Displays the maximum prefix length to be matched.

You can perform the following actions on entries in the Filtering table:

- Add—Opens the Add/Edit Filtering Entry dialog box for adding a new entry to the Filter table.
- Edit—Opens the Add/Edit Filtering Entry dialog box for modifying the selected filter.
- Delete—Removes the selected filter from the Filter table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Filtering Entry

Configuration > Routing > Dynamic Routing > OSPF > Filtering > Add/Edit Filtering Entry

The Add/Edit Filtering Entry dialog box lets you add new filters to the Filter table or to modify an existing filter. Some of the filter information cannot be changed when you edit an existing filter.

Fields

- OSPF Process—Choose the OSPF process associated with the filter entry. If you are editing an existing filter entry, you cannot modify this setting.
- Area ID—Choose the ID of the area associated with the filter entry. If you are editing an existing filter entry, you cannot modify this setting.
- Filtered Network—Enter the address and mask of the network being filtered using CIDR notation (a.b.c.d/m).

- Traffic Direction—Choose the traffic direction being filtered. Choose “Inbound” to filter LSAs coming into an OSPF area or “Outbound” to filter LSAs coming out of an OSPF area. If you are editing an existing filter entry, you cannot modify this setting.
- Sequence #—Enter a sequence number for the filter. Valid values range from 1 to 4294967294. When multiple filters apply to an LSA, the filter with the lowest sequence number is used.
- Action—Choose “Permit” to allow the LSA traffic or “Deny” to block the LSA traffic.
- Optional—Contains the optional settings for the filter.
  - Lower Range—Specify the minimum prefix length to be matched. The value of this setting must be greater than the length of the network mask entered in the Filtered Network field and less than or equal to the value, if present, entered in the Upper Range field.
  - Upper Range—Enter the maximum prefix length to be matched. The value of this setting must be greater than or equal to the value, if present, entered in the Lower Range field, or, if the Lower Range field is left blank, greater than the length of the network mask length entered in the Filtered Network field.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Interface

Configuration > Routing > Dynamic Routing > OSPF > Interface

The Interface pane lets you configure interface-specific OSPF authentication routing properties. For more information about configuring these properties, see the following:

- Interface > Authentication Tab
- Interface > Properties Tab

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Interface > Authentication Tab

Configuration > Routing > Dynamic Routing > OSPF > Interface > Authentication Tab

The Authentication tab displays the OSPF authentication information for the security appliance interfaces.

Fields

- Authentication Properties—Displays the authentication information for the security appliance interfaces. Double-clicking a row in the table opens the Edit OSPF Interface Properties dialog box for the selected interface.
  - Interface—Displays the interface name.
  - Authentication Type—Displays the type of OSPF authentication enabled on the interface. The authentication type can be one of the following values:

    None—OSPF authentication is disabled.

    Password—Clear text password authentication is enabled.

    MD5—MD5 authentication is enabled.

    Area—The authentication type specified for the area is enabled on the interface. Area authentication is the default value for interfaces. However, area authentication is disabled by default. So, unless you previously specified an area authentication type, interfaces showing Area authentication have authentication disabled.
- Edit—Opens the Edit OSPF Interface Properties dialog box for the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Edit OSPF Interface Authentication

Configuration > Routing > Dynamic Routing > OSPF > Interface > Authentication > Edit OSPF Interface Authentication

The Edit OSPF Interface Authentication dialog box lets you configure the OSPF authentication type and parameters for the selected interface.

Fields

- Interface—Displays the name of the interface for which authentication is being configured. You cannot edit this field.
- Authentication—Contains the OSPF authentication options.
  - None—Choose this option to disable OSPF authentication.
  - Password—Choose this option to use clear text password authentication. This is not recommended where security is a concern.
  - MD5—Choose this option to use MD5 authentication (recommended).
  - Area—(Default) Choose this option to use the authentication type specified for the area (see Add/Edit OSPF Area for information about configuring area authentication). Area authentication is disabled by default. So, unless you have previously specified an area authentication type, interfaces set to area authentication have authentication disabled until you configure area authentication.

- Authentication Password—Contains the settings for entering the password when password authentication is enabled.
  - Enter Password—Enter a text string of up to 8 characters.
  - Re-enter Password—Reenter the password.
- MD5 IDs and Keys—Contains the settings for entering the MD5 keys and parameters when MD5 authentication is enabled. All devices on the interface using OSPF authentication must use the same MD5 key and ID.
  - Enter MD5 ID and Key—Contains the settings for entering MD5 key information.

    Key ID—Enter a numerical key identifier. Valid values range from 1 to 255.

    Key—An alphanumeric character string of up to 16 bytes.
  - Add—Adds the specified MD5 key to the MD5 ID and Key table.
  - Delete—Removes the selected MD5 key and ID from the MD5 ID and Key table.
  - MD5 ID and Key—Displays the configured MD5 keys and key IDs.

    Key ID—Displays the key ID for the selected key.

    Key—Displays the key for the selected key ID.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Interface > Properties Tab

Configuration > Routing > Dynamic Routing > OSPF > Interface > Properties Tab

The Properties tab displays the OSPF properties defined for each interface in a table format.

Fields

- OSPF Interface Properties—Displays interface-specific OSPF properties. Double-clicking a row in the table opens the Edit OSPF Interface Properties dialog box for the selected interface.
  - Interface—Displays the name of the interface that the OSPF configuration applies to.
  - Broadcast—Displays “No” if the interface is set to non-broadcast (point-to-point). Displays “Yes” if the interface is set to broadcast. “Yes” is the default setting for Ethernet interfaces.
  - Cost—Displays the cost of sending a packet through the interface.
  - Priority—Displays the OSPF priority assigned to the interface.
  - MTU Ignore—Displays “No” if MTU mismatch detection is enabled. Displays “Yes” if the MTU mismatch detection is disabled.
  - Database Filter—Displays “Yes” if outgoing LSAs are filtered during synchronization and flooding. Displays “No” if filtering is not enabled.
- Edit—Opens the Edit OSPF Interface Properties dialog box for the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Edit OSPF Interface Properties

Configuration > Routing > Dynamic Routing > OSPF > Interface > Properties > Edit OSPF Interface Properties

Fields

- Interface—Displays the name of the interface for which you are configuring OSPF properties. You cannot edit this field.
- Broadcast—Check this check box to specify that the interface is a broadcast interface. This check box is selected by default for Ethernet interfaces. Uncheck this check box to designate the interface as a point-to-point, non-broadcast interface. Specifying an interface as point-to-point, non-broadcast lets you transmit OSPF routes over VPN tunnels.

  When an interface is configured as point-to-point, non-broadcast, the following restrictions apply:

  - You can define only one neighbor for the interface.
  - You need to manually configure the neighbor (see Static Neighbor).
  - You need to define a static route pointing to the crypto endpoint (see Static Routes).
  - If OSPF over the tunnel is running on the interface, regular OSPF with an upstream router cannot be run on the same interface.
  - You should bind the crypto-map to the interface before specifying the OSPF neighbor to ensure that the OSPF updates are passed through the VPN tunnel. If you bind the crypto-map to the interface after specifying the OSPF neighbor, use the **clear local-host all** command to clear OSPF connections so the OSPF adjacencies can be established over the VPN tunnel.
- Cost—Specify the cost of sending a packet through the interface. The default value is 10.
- Priority—Specify the OSPF router priority. When two routers connect to a network, both attempt to become the designated router. The devices with the higher router priority becomes the designated router. If there is a tie, the router with the higher router ID becomes the designated router.

  Valid values for this setting range from 0 to 255.The default value is 1. Entering 0 for this setting makes the router ineligible to become the designated router or backup designated router. This setting does not apply to interfaces that are configured as point-to-point non-broadcast interfaces.
- MTU Ignore—OSPF checks whether neighbors are using the same MTU on a common interface. This check is performed when neighbors exchange DBD packets. If the receiving MTU in the DBD packet is higher than the IP MTU configured on the incoming interface, OSPF adjacency will not be established.

- Database Filter—Check this check box to filter outgoing LSA interface during synchronization and flooding. By default, OSPF floods new LSAs over all interfaces in the same area, except the interface on which the LSA arrives. In a fully meshed topology, this can waste bandwidth and lead to excessive link and CPU usage. Checking this check box prevents flooding OSPF LSA on the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Edit OSPF Interface Advanced Properties

Configuration > Routing > Dynamic Routing > OSPF > Interface > Properties > Edit OSPF Interface Properties > Edit OSPF Interface Advanced Properties

The Edit OSPF Interface Advanced Properties dialog box lets you change the values for the OSPF hello interval, retransmit interval, transmit delay, and dead interval. Typically, you only need to change these values from the defaults if you are experiencing OSPF problems on your network.

Fields

- Hello Interval—Specifies the interval, in seconds, between hello packets sent on an interface. The smaller the hello interval, the faster topological changes are detected but the more traffic is sent on the interface. This value must be the same for all routers and access servers on a specific interface. Valid values range from 1 to 65535 seconds. The default value is 10 seconds.
- Retransmit Interval—Specifies the time, in seconds, between LSA retransmissions for adjacencies belonging to the interface. When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgement message. If the router receives no acknowledgement, it will resend the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links. Valid values range from 1 to 65535 seconds. The default value is 5 seconds.
- Transmit Delay—Specifies the estimated time, in seconds, required to send an LSA packet on the interface. LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links. Valid values range from 1 to 65535 seconds. The default value is 1 second.
- Dead Interval—Specifies the interval, in seconds, in which no hello packets are received, causing neighbors to declare a router down. Valid values range from 1 to 65535. The default value of this setting is four times the interval set by the Hello Interval field.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Redistribution

Configuration > Routing > Dynamic Routing > OSPF > Redistribution

The Redistribution pane displays the rules for redistributing routes from one routing domain to another.

Fields

The Redistribution table displays the following information. Double-clicking a table entry opens the Add/Edit OSPF Redistribution Entry dialog box for the selected entry.

- OSPF Process—Displays the OSPF process associated with the route redistribution entry.
- Protocol—Displays the source protocol the routes are being redistributed from. Valid entries are the following:
  - Static—The route is a static route.
  - Connected—The route was established automatically by virtue of having IP enabled on the interface. These routes are redistributed as external to the AS.
  - OSPF—The route is an OSPF route from another process.
- Match—Displays the conditions used for redistributing routes from one routing protocol to another.
- Subnets—Displays “Yes” if subnetted routes are redistributed. Does not display anything if only routes that are not subnetted are redistributed.
- Metric Value—Displays the metric that is used for the route. This column is blank for redistribution entries if the default metric is used.
- Metric Type—Displays “1” if the metric is a Type 1 external route, “2” if the metric is Type 2 external route.
- Tag Value—A 32-bit decimal value attached to each external route. This value is not used by OSPF itself. It may be used to communicate information between ASBRs. Valid values range from 0 to 4294967295.
- Route Map—Displays the name of the route map to apply to the redistribution entry.

You can perform the following actions on the Redistribution table entries:

- Add—Opens the Add/Edit OSPF Redistribution Entry dialog box for adding a new redistribution entry.
- Edit—Opens the Add/Edit OSPF Redistribution Entry dialog box for modifying the selected redistribution entry.
- Delete—Removes the selected redistribution entry from the Redistribution table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit OSPF Redistribution Entry

Configuration > Routing > Dynamic Routing > OSPF > Redistribution > Add/Edit OSPF Redistribution Entry

The Add/Edit OSPF Redistribution Entry dialog box lets you add a new redistribution rule to or edit an existing redistribution rule in the Redistribution table. Some of the redistribution rule information cannot be changed when you are editing an existing redistribution rule.

Fields

- OSPF Process—Choose the OSPF process associated with the route redistribution entry. If you are editing an existing redistribution rule, you cannot change this setting.
- Protocol—Choose the source protocol the routes are being redistributed from. You can choose one of the following options:
  - Static—The route is a static route.
  - Connected—The route was established automatically by virtue of having IP enabled on the interface. Connected routes are redistributed as external to the AS.
  - OSPF—The route is an OSPF route from another process.

    OSPF—Choose the OSPF process ID for the route being redistributed.
- Match—Displays the conditions used for redistributing routes from one routing protocol to another. The routes must match the selected condition to be redistributed. You can choose one or more of the following match conditions:
  - Internal—The route is internal to a specific AS.
  - External 1—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes.
  - External 2—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 external routes.
  - NSSA External 1—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes.
  - NSSA External 2—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes.
- Metric Value—Specify the metric value for the routes being redistributed. Valid values range from 1 to 16777214. When redistributing from one OSPF process to another OSPF process on the same device, the metric will be carried through from one process to the other if no metric value is specified. When redistributing other processes to an OSPF process, the default metric is 20 when no metric value is specified.
- Metric Type—Choose “1” if the metric is a Type 1 external route, “2” if the metric is a Type 2 external route.

- Tag Value—The tag value is a 32-bit decimal value attached to each external route. This is not used by OSPF itself. It may be used to communicate information between ASBRs. Valid values range from 0 to 4294967295.
- Use Subnets—Choose this check box to enable the redistribution of subnetted routes. Uncheck this check box to cause only routes that are not subnetted to be redistributed.
- Route Map—Enter the name of the route map to apply to the redistribution entry.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Static Neighbor

Configuration > Routing > Dynamic Routing > OSPF > Static Neighbor

The Static Neighbor pane displays manually defined neighbors; it does not display discovered neighbors.

You need to define a static neighbor for each point-to-point, non-broadcast interface. You also need to define a static route for each static neighbor in the Static Neighbor table.

Fields

- Static Neighbor—Displays information for the static neighbors defined for each OSPF process. Double-clicking a row in the table opens the Add/Edit OSPF Neighbor Entry dialog box.
  - OSPF Process—Displays the OSPF process associated with the static neighbor.
  - Neighbor—Displays the IP address of the static neighbor.
  - Interface—Displays the interface associated with the static neighbor.
- Add—Opens the Add/Edit OSPF Neighbor Entry dialog box. Use this button to define a new static neighbor.
- Edit—Opens the Add/Edit OSPF Neighbor Entry dialog box. Use this button to change the settings for a static neighbor.
- Delete—Removes the selected entry from the Static Neighbor table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit OSPF Neighbor Entry

Configuration > Routing > Dynamic Routing > OSPF > Static Neighbor > Add/Edit OSPF Neighbor Entry

The Add/Edit OSPF Neighbor Entry dialog box lets you define a new static neighbor or change information for an existing static neighbor.

You must define a static neighbor for each point-to-point, non-broadcast interface.

Restrictions

- You cannot define the same static neighbor for two different OSPF processes.
- You need to define a static route for each static neighbor (see Static Routes).

Fields

- OSPF Process—Choose the OSPF process associated with the static neighbor. If you are editing an existing static neighbor, you cannot change this value.
- Neighbor—Enter the IP address of the static neighbor.
- Interface—Choose the interface associated with the static neighbor. If you are editing an existing static neighbor, you cannot change this value.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Summary Address

Configuration > Routing > Dynamic Routing > OSPF > Summary Address

The Summary Address pane displays information about the summary addresses configured for each OSPF routing process.

Routes learned from other routing protocols can be summarized. The metric used to advertise the summary is the smallest metric of all the more specific routes. Summary routes help reduce the size of the routing table.

Using summary routes for OSPF causes an OSPF ASBR to advertise one external route as an aggregate for all redistributed routes that are covered by the address. Only routes from other routing protocols that are being redistributed into OSPF can be summarized.

Fields

The following information appears in the Summary Address table. Double-clicking an entry in the table opens the Add/Edit OSPF Summary Address Entry dialog box for the selected entry.

- OSPF Process—Displays the OSPF process associated with the summary address.
- IP Address—Displays the IP address of the summary address.
- Netmask—Displays the network mask of the summary address.

- Advertise—Displays “Yes” if the summary routes are advertised. Displays “No” if the summary route is not advertised.
- Tag—Displays a 32-bit decimal value attached to each external route. This value is not used by OSPF itself. It may be used to communicate information between ASBRs.

You can perform the following actions on the entries in the Summary Address table:

- Add—Opens the Add/Edit OSPF Summary Address Entry dialog box for adding new summary address entries.
- Edit—Opens the Add/Edit OSPF Summary Address Entry dialog box for editing the selected entry.
- Delete—Removes the selected summary address entry from the Summary Address table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit OSPF Summary Address Entry

Configuration > Routing > Dynamic Routing > OSPF > Summary Address > Add/Edit OSPF Summary Address Entry

The Add/Edit OSPF Summary Address Entry dialog box lets you add new entries to or modify existing entries in the Summary Address table. Some of the summary address information cannot be changed when editing an existing entry.

Fields

- OSPF Process—Choose the OSPF process associated with the summary address. You cannot change this information when editing an existing entry.
- IP Address—Enter the IP address of the summary address. You cannot change this information when editing an existing entry.
- Netmask—Enter the network mask for the summary address, or choose the network mask from the list of common masks. You cannot change this information when editing an existing entry.
- Advertise—Check this check box to advertise the summary route. Uncheck this check box to suppress routes that fall under the summary address. By default this check box is selected.
- Tag—(Optional) The tag value is a 32-bit decimal value attached to each external route. This is not used by OSPF itself. It may be used to communicate information between ASBRs. Valid values range from 0 to 4294967295.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Virtual Link

Configuration > Routing > Dynamic Routing > OSPF > Virtual Link

If you add an area to an OSPF network, and it is not possible to connect the area directly to the backbone area, you need to create a virtual link. A virtual link connects two OSPF devices that have a common area, called the transit area. One of the OSPF devices must be connected to the backbone area.

Fields

The Virtual Link table displays the following information. Doubling-clicking an entry in the table opens the Add/Edit Virtual Link dialog box for the selected entry.

- OSPF Process—Displays the OSPF process associated with the virtual link.
- Area ID—Displays the ID of the transit area.
- Peer Router ID—Displays the router ID of the virtual link neighbor.
- Authentication—Displays the type of authentication used by the virtual link:
  - None—No authentication is used.
  - Password—Clear text password authentication is used.
  - MD5—MD5 authentication is used.

You can perform the following actions on the entries in the Virtual Link table:

- Add—Opens the Add/Edit Virtual Link dialog box for adding a new entry to the Virtual Link table.
- Edit—Opens the Add/Edit Virtual Link dialog box for the selected entry.
- Delete—Removes the selected entry from the Virtual Link table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Virtual Link

Configuration > Routing > Dynamic Routing > OSPF > Virtual Link > Add/Edit Virtual Link

The Add/Edit Virtual Link dialog box lets you define new virtual links or change the properties of existing virtual links.

Fields

- OSPF Process—Choose the OSPF process associated with the virtual link. If you are editing an existing virtual link, you cannot change this value.
- Area ID—Choose the area shared by the neighbor OSPF devices. The selected area cannot be an NSSA or a Stub area. If you are editing an existing virtual link, you cannot change this value.
- Peer Router ID—Enter the router ID of the virtual link neighbor. If you are editing an existing virtual link, you cannot change this value.
- Advanced—Opens the Advanced OSPF Virtual Link Properties dialog box. You can configure the OSPF properties for the virtual link in this area. These properties include authentication and packet interval settings.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Advanced OSPF Virtual Link Properties

Configuration > Routing > Dynamic Routing > OSPF > Virtual Link > Add/Edit Virtual Link > Advanced OSPF Virtual Link Properties

The Advanced OSPF Virtual Link Properties dialog box lets you configure OSPF authentication and packet intervals.

Fields

- Authentication—Contains the OSPF authentication options.
  - None—Choose this option to disable OSPF authentication.
  - Password—Choose this option to use clear text password authentication. This is not recommended where security is a concern.
  - MD5—Choose this option to use MD5 authentication (recommended).
- Authentication Password—Contains the settings for entering the password when password authentication is enabled.
  - Enter Password—Enter a text string of up to 8 characters.
  - Re-enter Password—Reenter the password.
- MD5 IDs and Keys—Contains the settings for entering the MD5 keys and parameters when MD5 authentication is enabled. All devices on the interface using OSPF authentication must use the same MD5 key and ID.
  - Enter MD5 ID and Key—Contains the settings for entering MD5 key information.

    Key ID—Enter a numerical key identifier. Valid values range from 1 to 255.

    Key—An alphanumeric character string of up to 16 bytes.
  - Add—Adds the specified MD5 key to the MD5 ID and Key table.

  - Delete—Removes the selected MD5 key and ID from the MD5 ID and Key table.
  - MD5 ID and Key—Displays the configured MD5 keys and key IDs.

    Key ID—Displays the key ID for the selected key.

    Key—Displays the key for the selected key ID.

- Intervals—Contains the settings for modifying packet interval timing.
  - Hello Interval—Specifies the interval, in seconds, between hello packets sent on an interface. The smaller the hello interval, the faster topological changes are detected but the more traffic is sent on the interface. This value must be the same for all routers and access servers on a specific interface. Valid values range from 1 to 65535 seconds. The default value is 10 seconds.
  - Retransmit Interval—Specifies the time, in seconds, between LSA retransmissions for adjacencies belonging to the interface. When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgement message. If the router receives no acknowledgement, it will resend the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links. Valid values range from 1 to 65535 seconds. The default value is 5 seconds.
  - Transmit Delay—Specifies the estimated time, in seconds, required to send an LSA packet on the interface. LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links. Valid values range from 1 to 65535 seconds. The default value is 1 second.
  - Dead Interval—Specifies the interval, in seconds, in which no hello packets are received, causing neighbors to declare a router down. Valid values range from 1 to 65535. The default value of this field is four times the interval set by the Hello Interval field.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

RIP

Configuration > Routing > Dynamic Routing > RIP

RIP is a distance-vector routing protocol that uses hop count as the metric for path selection. When RIP is enabled on an interface, the interface exchanges RIP broadcasts with neighboring devices to dynamically learn about and advertise routes.

The security appliance support both RIP version 1 and RIP version 2. RIP version 1 does not send the subnet mask with the routing update. RIP version 2 sends the subnet mask with the routing update and supports variable-length subnet masks. Additionally, RIP version 2 supports neighbor authentication when routing updates are exchanged. This authentication ensures that the security appliance receives reliable routing information from a trusted source.

You can have two OSPF routing processes and one RIP routing process running on the security appliance at the same time.

Limitations

RIP has the following limitations:

- The security appliance cannot pass RIP updates between interfaces.
- RIP Version 1 does not support variable-length subnet masks.
- RIP has a maximum hop count of 15. A route with a hop count greater than 15 is considered unreachable.
- RIP convergence is relatively slow compared to other routing protocols.
- You can only enable a single RIP process on the security appliance.

RIP Version 2 Notes

The following information applies to RIP Version 2 only:

- If using neighbor authentication, the authentication key and key ID must be the same on all neighbor devices that provide RIP version 2 updates to the interface.
- With RIP version 2, the security appliance transmits and receives default route updates using the multicast address 224.0.0.9. In passive mode, it receives route updates at that address.
- When RIP version 2 is configured on an interface, the multicast address 224.0.0.9 is registered on that interface. When a RIP version 2 configuration is removed from an interface, that multicast address is unregistered.

Global Setup

Configuration > Routing > Dynamic Routing > RIP >Global Setup

Use the Global Setup pane to enable RIP on the security appliance and to configure global RIP protocol parameters. You can only enable a single RIP process on the security appliance.

Fields

- Enable RIP Routing—Check this check box to enable RIP routing on the security appliance. When you enable RIP, it is enabled on all interfaces. Checking this check box also enables the other fields on this pane. Uncheck this check box to disable RIP routing on the security appliance.
- Enable Auto-summarization—Clear this check box to disable automatic route summarization. Check this check box to reenable automatic route summarization. RIP Version 1 always uses automatic summarization. You cannot disable automatic summarization for RIP Version 1. If you are using RIP Version 2, you can turn off automatic summarization by unchecking this check box. Disable automatic summarization if you must perform routing between disconnected subnets. When automatic summarization is disabled, subnets are advertised.
- Enable RIP version—Check this check box to specify the vesion of RIP used by the security appliance. If this check boxis cleared, then the security appliance sends RIP Version 1 updates and accepts RIP Version 1 & Version 2 updates. This setting can be overridden on a per-interface basis in the Interface pane.
  - Version 1—Specifies that the security appliance only sends and receives RIP Version 1 updates. Any version 2 updates received are dropped.
  - Version 2—Specifies that the security appliance only sends and receives RIP Version 2 updates. Any version 1 updates received are dropped.

- Enable default information originate—Check this check box to generate a default route into the RIP routing process. You can configure a route map that must be satisfied before the default route can be generated.
  - Route-map—Enter the name of the route map to apply. The routing process generates the default route if the route map is satisfied.
- IP Network to Add—Defines a network for the RIP routing process. The network number specified must not contain any subnet information. There is no limit to the number of network you can add to the security appliance configuration. RIP routing updates will be sent and received only through interfaces on the specified networks. Also, if the network of an interface is not specified, the interface will not be advertised in any RIP updates.
  - Add—Click this button to add the specified network to the list of networks.
  - Delete—Click this button to removed the selected network from the list of networks.
- Configure interfaces as passive globally—Check this check box to set all interfaces on the security appliance to passive RIP mode. The security appliance listens for RIP routing broadcasts on all interfaces and uses that information to populate the routing tables but do not broadcast routing updates. To set specific interfaces to passive RIP, use the Passive Interfaces table.
- Passive Interfaces table—Lists the configured interfaces on the security appliance. Check the check box in the Passive column for those interfaces you want to operate in passive mode. The other interfaces will still send and receive RIP broadcasts.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Interface

Configuration > Routing > Dynamic Routing > RIP > Interface

The Interface pane allows you to configure interface-specific RIP settings, such as the version of RIP the interface sends and receives and the authentication method, if any, used for the RIP broadcasts.

Fields

- Interface table—(*Display only*) Each row displays the interface-specific RIP settings for an interface. Double-clicking a row for that entry opens the Edit RIP Interface Entry dialog box for that interface.
- Edit—Opens the Edit RIP Interface Entry dialog box for the interface selected in the Interface table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Edit RIP Interface Entry

Configuration > Routing > Dynamic Routing > RIP > Interface > Edit RIP Interface Entry

The Edit RIP Interface Entry dialog box allows you to configurethe interface-specific RIP settings.

Fields

- Override Global Send Version—Check this check box to specify the RIP version sent by the interface. You can select the following options:
  - Version 1
  - Version 2
  - Version 1 & 2

  Unchecking this check box restores the global setting.
- Override Global Receive Version—Check this check box to specify the RIP version accepted by the interface. If a RIP updated from an unsuppored version of RIP is received by the interface, it is dropped. You can select the following options:
  - Version 1
  - Version 2
  - Version 1 & 2

  Unchecking this check box restores the global setting.
- Enable Authentication—Check this check box to enable RIP authentication. Uncheck this check box to disable RIP broadcast authentication.
  - Key—The key used by the authentication method. Can conting up to 16 characters.
  - Key ID—The key ID. Valid values are from 0 to 255.
  - Authentication Mode—You can select the following authentication modes:

    MD5—Uses MD5 for RIP message authentication.

    Text—Uses cleartext for RIP message authentication (not recommended).

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Filter Rules

Configuration > Routing > Dynamic Routing > RIP > Filter Rules

Filter rules allow you to filter the network received in RIP routing updates or sent in RIP routing updates. Each filter rule consists of one or more network rules.

Fields

- Filter Rules table—Displays the configured RIP filter rules.
- Add—Clicling this button opens the Add/Edit Filter Rule dialog box. The new filter rule is added to the bottom of the list.
- Edit—Clicking this button opens the Add/Edit Filter Rule dialog box for the selected filter rule.
- Delete—Clicking this button deletes the selected filter rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Filter Rule

Configuration > Routing > Dynamic Routing > RIP > Filter Rules > Add/Edit Fitler Rule

Use the Add/Edit Filter Rule pane to create filter rules. You can create filter rules that apply to all interfaces or that apply to a specific interface.

Fields

- Direction—Select one of the following directions for the filter to act upon:
  - In—Filters networks on incoming RIP updates.
  - Out—Filters networks from outgoing RIP updates.
- Interface—You can select a specific interface for the filter rule, or you can select the All Interfaces option to apply the filter to all interfaces.
- Action—(*Display only*) Displays Permit if the specified network is not filtered from incoming or outgoing RIP advertisements. Displays Deny if the specified network is to be filtered from incoming or outgoing RIP advertisements.
- IP Address—(*Display only*) Displays the IP address of the network being filtered.
- Netmask—(*Display only*) Displays the network mask applied to the IP address.
- Insert—Click this button to add a network rule above the selected rule in the list. Clicking this button opens the Network Rule dialog box.
- Edit—Click this button to edit the selected rule. Clicking this button opens the Network Rule dialog box.
- Add—Click this button to add a network rule below the selected rule in the list. Clicking this button opens the Network Rule dialog box.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Network Rule

Configuration > Routing > Dynamic Routing > RIP > Filter Rules > Add/Edit Fitler Rule > Network Rule

The Network Rule pane allows you to configure permit and deny rules for specific networks in a filter rule.

Fields

- Action—Select Permit to allow the specified network to be advertised in RIP updates or accepted into the RIP routing process. Select Deny to prevent the specified network from being advertised in RIP updates or accepted into the RIP routing process.
- IP Address—Type IP address of the network being permitted or denied.
- Netmask—Specify the network mask applied to the network IP address. You can type a network mask into this field or select one of the common masks from the list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Route Redistribution

Configuration > Routing > Dynamic Routing > RIP > Route Redistribution

The Route Redistribution pane displays the routes that are being redistributed from other routing processes into the RIP routing process.

Fields

- Protocol—(*Display only*) Displays the routing protocol being redistributed into the RIP routing process:
  - Static—Static routes.
  - Connected—Directly connected networks.
  - OSPF—Networks discovered by the specified OSPF routing process.
- Metric—The RIP metric being applied to the redistributed routes.

- Match—(*Display only*) Displays the type of OSPF routes being redistributed into the RIP routing process. If the Match column is blank for an OSPF redistribution rule, Internal, External 1, and External 2 routes are redistributed into the RIP routing process.
- Route Map—(*Display only*) Displays the name of the route map, if any, being applied to the redistribution. Route maps are used to specify with greater detail which routes from the specified routing process are redistributed into RIP.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Route Redistribution

Configuration > Routing > Dynamic Routing > RIP > Route Redistribution > Add/Edit Route Redistribution

Use the Add Route Redistribution dialog box to add a new redistribution rule. Use the Edit Route Redistribution dialog box to change an existing rule.

Fields

- Protocol—Choose the routing protocol to redistribute into the RIP routing process:
  - Static—Static routes.
  - Connected—Directly connected networks.
  - OSPF and OSPF ID—Routes discovered by the OSPF routing process. If you choose OSPF, you must also enter the OSPF process ID. Additionally, you can select the specific types of OSPF routes to redistribute from the Match area.
- Route Map—Specifies the name of a route map that must be satisfied before the route can be redistributed into the RIP routing process.
- Configure Metric Type—Check this checkbox to specify a metric for the redistributed routes. If not specified, the routes are assigned a metric of 0.
  - Transparent—Choose this option
  - Value—Choose this to assign a specific metric value. You can enter a value from 0 to 16.
- Match—If you are redistributing OSPF routes into the RIP routing process, you can choose specific types of OSPF routes to redistribute by checking the check box next to the route type. If you do not check any route types, Internal, External 1, and External 2 routes are redistributed by default.
  - Internal—Routes internal to the AS are redistributed.
  - External 1—Type 1 routes external to the AS are redistribued.
  - External 2—Type 2 routes external to the AS are redistributed.
  - NSSA External 1—Type 1 routes external to an NSSA are redistributed.
  - NSSA External 2—Type 2 routes external to an NSSA are redistributed.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Static Routes

Multiple context mode does not support dynamic routing, so you must define static routes for any networks to which the security appliance is not directly connected.

In transparent firewall mode, for traffic that originates on the security appliance and is destined for a non-directly connected network, you need to configure either a default route or static routes so the security appliance knows out of which interface to send traffic. Traffic that originates on the security appliance might include communications to a syslog server, Websense or N2H2 server, or AAA server. If you have servers that cannot all be reached through a single default route, then you must configure static routes.

The simplest option is to configure a default route to send all traffic to an upstream router, relying on the router to route the traffic for you. However, in some cases the default gateway might not be able to reach the destination network, so you must also configure more specific static routes. For example, if the default gateway is on the outside interface, the default route cannot direct traffic to any inside networks that are not directly connected to the security appliance.

You can also use static route in conjunction with dynamic routing protocols to provide a floating static route that is used when the dynamically discovered route goes down. If you create a static route with an administrative distance greater than the administrative distance of the dynamic routing protocol, then a route to the specified destination discovered by the routing protocol takes precedence over the static route. The static route is used only if the dynamically discovered route is removed from the routing table.

Static routes remain in the routing table even if the specified gateway becomes unavailable (see Static Route Tracking, page 14-30, for the exception to this). If the specified gateway becomes unavailable, you need to remove the static route from the routing table manually. However, static routes are removed from the routing table if the associated interface on the security appliance goes down. They are reinstated when the interface comes back up.

You can define up to three equal cost routes to the same destination per interface. ECMP is not supported across multiple interfaces. With ECMP, the traffic is not necessarily divided evenly between the routes; traffic is distributed among the specified gateways based on an algorithm that hashes the source and destination IP addresses.

The default route identifies the gateway IP address to which the security appliance sends all IP packets for which it does not have a learned or static route. A default route is simply a static route with 0.0.0.0/0 as the destination IP address. Routes that identify a specific destination take precedence over the default route.

You can define up to three equal cost default route entries per device. Defining more than one equal cost default route entry causes the traffic sent to the default route to be distributed among the specified gateways. When defining more than one default route, you must specify the same interface for each entry.

If you attempt to define more than three equal cost default routes, or if you attempt to define a default route with a different interface than a previously defined default route, youwill receive an error message.

You can define a separate default route for tunneled traffic along with the standard default route. When you create a default route with the **tunneled** option, all encrypted traffic that arrives on the security appliance and that cannot be routed using learned or static routes is sent to this route. Otherwise, if the traffic is not encrypted, the standard default route entry is used. You cannot define more than one default route with the **tunneled** option; ECMP for tunneled traffic is not supported.

For more information about viewing and configuring static and default routes with ASDM, see Field Information for Static Routes, page 14-31.

Static Route Tracking

It is not always possible to use dynamic routing protocols on the security appliance, such as when the security appliance is in multiple context mode or transparent mode. In these cases, you must use static routes.

One of the problems with static routes is that there is no inherent mechanism for determining if the route is up or down. They remain in the routing table even if the next hop gateway goes down. They are only removed from the routing table if the associated interface on the security appliance goes down.

The static route tracking feature provides a method for tracking the availability of a static route and installing a backup route if the primary route should fail. This allows you to, for example, define a default route to an ISP gateway and a backup default route to a secondary ISP in case the primary ISP becomes unavailable.

The security appliance does this by associating a static route with a monitoring target that you define. It monitors the target using ICMP echo requests. If an echo reply is not received within a specified time period, the object is considered down and the associated route is removed from the routing table. A previously configured backup route is used in place of the removed route.

When selecting a monitoring target, you need to make sure that it can respond to ICMP echo requests. The target can be any network object that responds to ICMP echo requests. Consider choosing:

- the ISP gateway (for dual ISP support) address
- the next hop gateway address (if you are concerned about the availability of the gateway)
- a server, such as a AAA server, that the security appliance needs to communicate with
- a persistent network object on the destination network (a desktop or notebook computer that may be shut down at night is not a good choice)

For more information about configuring static route tracking, see Configuring Static Route Tracking, page 14-30. To monitor the static route tracking process, see interface connection, page 40-9.

Configuring Static Route Tracking

This procedure provides an overview of configuring static route tracking. For specific information about the various fields used to configure this feature, see Field Information for Static Routes, page 14-31

To configure tracking for a static route, perform these steps:

Step 1 Choose a target of interest. Make sure the target responds to echo requests.

Step 2 Open the Static Routes page. Go to **Configuration > Routing > Static Routes**.

Step 3 Click **Add** to configure a static route that is to be used based on the availability of your selected target of interest. You must enter the Interface, IP Address, Mask, Gateway, and Metric for this route. See Add/Edit Static Route, page 14-32, for more information about these fields.

Step 4 Choose **Tracked** in the Options area for this route.

Step 5 Configure the tracking properties. You must enter a unique Track ID, a unique SLA ID, and the IP address of your target of interest. See Add/Edit Static Route, page 14-32, for more information about these fields.

Step 6 (Optional) To configure the monitoring properties, click **Monitoring Options** in the Add Static Route dialog box. See Route Monitoring Options, page 14-33, for more information about the monitoring properties.

Step 7 Click **OK** to save your changes.

The monitoring process begins as soon as you save the tracked route.

Step 8 Create a secondary route. The secondary route is a static route to the same destination as the tracked route, but through a different interface or gateway. You must assign this route a higher administrative distance (metric) than your tracked route.

Field Information for Static Routes

For information about a specific pane, see the following topics:

- Static Routes, page 14-31
- Add/Edit Static Route, page 14-32
- Route Monitoring Options, page 14-33

Static Routes

Configuration > Routing > Dynamic Routing > Static Routes

The Static Route pane lets you create static routes that will access networks connected to a router on any interface. To enter a default route, set the IP address and mask to 0.0.0.0, or the shortened form of 0.

If an IP address from one security appliance interface is used as the gateway IP address, the security appliance will ARP the designated IP address in the packet instead of ARPing the gateway IP address.

Leave the Metric at the default of 1 unless you are sure of the number of hops to the gateway router.

Fields

The Static Route pane shows the Static Route table:

- Interface—(*Display only*) Lists the internal or external network interface name enabled in Interfaces.
- IP Address—(*Display only*) Lists the internal or external network IP address. Use **0.0.0.0** to specify a default route. The **0.0.0.0** IP address can be abbreviated as **0**.
- Netmask—(*Display only*) Lists the network mask address that applies to the IP address. Use **0.0.0.0** to specify a default route. The **0.0.0.0** netmask can be abbreviated as **0**.
- Gateway IP—(*Display only*) Lists the IP address of the gateway router, which is the next hop address for this route.

- Metric—(*Display only*) Lists the administrative distance of the route. The default is 1 if a metric is not specified.
- Options—(*Display only*) Displays any options specified for thestatic route.
  - None—No options are specified for the static route.
  - Tunneled—Specifies route as the default tunnel gateway for VPN traffic. Used only for default route. You can only configure one tunneled route per device. The tunneled option is not supported under transparent mode.
  - Tracked—Specifies that the route is tracked. The tracking object ID and the address of the tracking target are also displayed. The tracked option is only supported in single, routed mode.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Static Route

Configuration >Routing > Static Routes > Add/Edit Static Route

Startup Wizard > Static Routes >Add/Edit Static Route

Use the Add/Edit Static Route dialog box to configure the static route properties. This dialog box is availble from both the Static Routes screen in the Startup Wizard and the Configuration > Routing > Static Route pane.

Fields

- Interface Name—Select the egress interface for the route.
- IP Address—Specifies the internal or external network IP address. Use **0.0.0.0** to specify a default route. The **0.0.0.0** IP address can be abbreviated as **0**.
- Mask—Specifies the network mask address that applies to the IP address. Use **0.0.0.0** to specify a default route. The **0.0.0.0** netmask can be abbreviated as **0**.
- Gateway IP—Specifies the IP address of the gateway router, which is the next hop address for this router.
- Metric—Lets you specify the the administrative distance of the route. The default is **1** if a metric is not specified.

The following options are available for static routes. You can select only one of these options for a static route. By default, no option (None) is selected.

- None—No options are specified for the static route.
- Tunneled—Used only for default route. Only one default tunneled gateway is allowed per security appliance. Tunneled option is not supported under transparent mode.
- Tracked—Select this option to specify that the route is tracked. Specifying this option starts the route tracking process.
  - Track ID—A unique identifier for the route tracking process.

- Track IP Address/DNS Name—Enter the IP address or hostname of the target being tracked. Typically, this would be the IP address of the next hop gateway for the route, but it could be any network object available off of that interface.
- SLA ID—A unique identifier for the SLA monitoring process.
- Monitor Options—Click this button to open the Route Monitoring Options dialog box. In the Route Monitoring Options dialog box you can configure the parameters of the tracked object monitoring process.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Route Monitoring Options

Configuration > Routing > Static Routes > Add/Edit Static Route > Route Monitoring Options

Use the Route Monitoring Options dialog box to change the tracking object monitoring properties.

Fields

- Frequency—Enter how often, in seconds, the security appliance should test for the presence of the tracking target. The default value is 60 seconds. Valid values are from 1 to 604800 seconds.
- Threshold—Enter the amount of time, in milliseconds, that indicates an over-threshold event. This value cannot be more than the timeout value.
- Timeout—Enter the amount of time, in milliseconds, the route monitoring operation should wait for a response from the request packets. The default value is 5000 milliseconds. Valid values are from 0 to 604800000 milliseconds.
- Data Size—Enter the size of data payload to use in the echo request packets. The default value is 28. Valid values are from 0 to 16384.

  Note This setting specifies the size of the payload only; it does not specify the size of the entire packet.

- ToS—Enter a value for the type of service byte in the IP header of the echo request. The default value is 0. Valid values are from 0 to 255.
- Number of Packets—The number of echo requests to send for each test. The default value is 1. Valid values are from 1 to 100.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

ASR Group

Configuration > Routing > ASR Group

Use the ASR Group screen to assign asynchronous routing group ID numbers to interfaces.

In some situations, return traffic for a session may be routed through a different interface than it originated from. In failover configurations, return traffic for a connection that originated on one unit may return through the peer unit. This most commonly occurs when two interfaces on a single security appliance, or two security appliances in a failover pair, are connected to different service providers and the outbound connection does not use a NAT address. By default, the security appliance drops the return traffic because there is no connection information for the traffic.

You can prevent the return traffic from being dropped using an ASR Group on interfaces where this is likely to occur. When an interface configured with an ASR Group receives a packet for which it has no session information, it checks the session information for the other interfaces that are in the same group.

Note You must enable Stateful Failover for session information to be passed from the standby failover group to the active failover group.

If it does not find a match, the packet is dropped. If it finds a match, then one of the following actions occurs:

- If the incoming traffic originated on a peer unit in a failover configuration, some or all of the layer 2 header is rewritten and the packet is redirected to the other unit. This redirection continues as long as the session is active.
- If the incoming traffic originated on a different interface on the same unit, some or all of the layer 2 header is rewritten and the packet is reinjected into the stream.

Fields

The **ASR Group** table displays the following information for each interface on the security appliance:

- **Interface**—Displays the name of the interface on the security appliance.
- **ASR Group ID**—Displays the number of the ASR Group the interface belongs to. If the interface has not been assigned an ASR Group number, this column displays "-- None --". Valid values are from 1 to 32.

  To assign an ASR Group number to an interface, click the **ASR Group ID** cell in the row of the desired interface. A list of valid ASR Group number appears. Select the desired ASR Group number from the list. You can assign a maximum of 8 interfaces to a single ASR Group. If other contexts have interfaces assigned to an ASR Group, those interface count against the total of 8, even for the context currently being configured.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	•	—

Proxy ARPs

Configuration > Routing > Proxy ARPs

In rare circumstances, you might want to disable proxy ARP for global addresses.

When a host sends IP traffic to another device on the same Ethernet network, the host needs to know the MAC address of the device. ARP is a Layer 2 protocol that resolves an IP address to a MAC address. A host sends an ARP request asking "Who is this IP address?" The device owning the IP address replies, "I own that IP address; here is my MAC address."

Proxy ARP is when a device responds to an ARP request with its own MAC address, even though the device does not own the IP address. The security appliance uses proxy ARP when you configure NAT and specify a global address that is on the same network as the security appliance interface. The only way traffic can reach the hosts is if the security appliance uses proxy ARP to claim that the security appliance MAC address is assigned to destination global addresses.

Fields

- Interface—Lists the interface names.
- Proxy ARP Enabled—Shows whether proxy ARP is enabled or disabled for NAT global addresses, Yes or No.
- Enable—Enables proxy ARP for the selected interface. By default, proxy ARP is enabled for all interfaces.
- Disable—Disables proxy ARP for the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Configuring Multicast Routing

Multicast routing is supported in single, routed mode only. This section contains the following topics:

- Multicast, page 15-1—enable or disable multicast routing on the security appliance.
- IGMP, page 15-2—configure IGMP on the security appliance.
- Multicast Route, page 15-7—define static multicast routes.
- MBoundary, page 15-8—configure boundaries for administratively-scoped multicast addresses.
- MForwarding, page 15-10—enable or disable multicast forwarding on a per-interface basis.
- PIM, page 15-11—configure PIM on the security appliance.

Multicast

Configuration > Routing > Multicast

The Multicast pane lets you enable multicast routing on the security appliance. Enabling multicast routing enables IGMP and PIM on all interfaces by default. IGMP is used to learn whether members of a group are present on directly attached subnets. Hosts join multicast groups by sending IGMP report messages. PIM is used to maintain forwarding tables to forward multicast datagrams.

Fields

Enable Multicast Routing—Check this check box to enable IP multicast routing on the security appliance. Uncheck this check box to disable IP multicast routing. By default, multicast is disabled. Enabling multicast enables multicast on all interfaces. You can disable multicast on a per-interface basis.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

IGMP

IP hosts use IGMP to report their group memberships to directly connected multicast routers. IGMP uses group address (Class D IP addresses). Host group addresses can be in the range 224.0.0.0 to 239.255.255.255. The address 224.0.0.0 is never assigned to any group. The address 224.0.0.1 is assigned to all systems on a subnet. The address 224.0.0.2 is assigned to all routers on a subnet.

For more information about configuring IGMP on the security appliance, see the following:

- Access Group
- Join Group
- Protocol
- Static Group

Access Group

Configuration > Routing > Multicast > IGMP > Access Group

Access groups control the multicast groups that are allowed on an interface.

Fields

- Access Groups—Displays the access groups defined for each interface.

  The table entries are processed from the top down. Place more specific entries near the top of the table and more generic entries further down. For example, place an access group entry that permits a specific multicast group near the top of the table and an access group entry below that denies a range of multicast groups, including the group in the permit rule. The group is permitted because the permit rule is enforced before the deny rule.

  Double-clicking an entry in the table opens the Add/Edit Access Group dialog box for the selected entry.

  - Interface—Displays the interface the access group is associated with.
  - Action—Displays “Permit” if the multicast group address is permitted by the access rule. Displays “Deny” if the multicast group address is denied by the access rule.
  - Multicast Group Address—Displays the multicast group address that the access rule applies to.
  - Netmask—Displays the network mask for the multicast group address.
- Insert Before—Opens the Add/Edit Access Group dialog box. Use this button to add a new access group entry before the selected entry in the table.
- Insert After—Opens the Add/Edit Access Group dialog box. Use this button to add a new access group entry after the selected entry in the table.
- Add—Opens the Add/Edit Access Group dialog box. Use this button to add a new access group entry at the bottom of the table.
- Edit—Opens the Add/Edit Access Group dialog box. Use this button to change the information for the selected access group entry.
- Delete—Removes the selected access group entry from the table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Access Group

Configuration > Routing > Multicast > IGMP > Access Group > Add/Edit Access Group

The Add Access Group dialog box lets you add a new access group to the Access Group Table. The Edit Access Group dialog box lets you change information for an existing access group entry. Some fields may be locked when editing existing entries.

Fields

- Interface—Choose the interface the access group is associated with. You cannot change the associated interface when you are editing an existing access group.
- Action—Choose "permit" to allow the multicast group on the selected interface. Choose "deny" to filter the multicast group from the selected interface.
- Multicast Group Address—Enter the address of the multicast group the access group applies to.
- Netmask—Enter the network mask for the multicast group address or choose one of the common network masks from the list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Join Group

Configuration > Routing > Multicast > IGMP > Join Group

You can configure the security appliance to be a member of a multicast group. The Join Group pane displays the multicast groups the security appliance is a member of.

Note If you simply want to forward multicast packets for a specific group to an interface without the security appliance accepting those packets as part of the group, see Static Group.

Fields

- Join Group—Displays the multicast group membership for each interface.
  - Interface—Displays the name of the security appliance interface.

  - Multicast Group Address—Displays the address of a multicast group that the interface belongs to.
- Add—Opens the Add/Edit IGMP Join Group dialog box. Use this button to add a new multicast group membership to an interface.
- Edit—Opens the Add/Edit IGMP Join Group dialog box. Use this button to edit an existing multicast group membership entry.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit IGMP Join Group

Configuration > Routing > Multicast > IGMP > Join Group > Add/Edit IGMP Join Group

Use the Add IGMP Join Group dialog box to configure an interface to be a member of a multicast group. Use the Edit IGMP Join Group dialog box to change existing membership information.

Fields

- Interface—Choose the name of the security appliance interface that you are configuring multicast group membership for. If you are editing an existing entry, you cannot change this value.
- Multicast Group Address—Enter the address of a multicast group in this field. The group address must be from 224.0.0.0 to 239.255.255.255.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Protocol

Configuration > Routing > Multicast > IGMP > Protocol

The Protocol pane displays the IGMP parameters for each interface on the security appliance.

Fields

- Protocol—Displays the IGMP parameters set on each interface. Double-clicking a row in the table opens the Configure IGMP Parameters dialog box for the selected interface.

- Interface—Displays the name of the interface.
- Enabled—Displays “Yes” if IGMP is enabled on the interface. Displays “No” if IGMP is disabled on the interface.
- Version—Displays the version of IGMP enabled on the interface.
- Query Interval—Displays the interval, in seconds, at which the designated router sends IGMP host-query messages.
- Query Timeout—Displays the period of time before which the security appliance takes over as the querier for the interface after the previous querier has stopped doing so.
- Response Time—Displays the maximum response time, in seconds, advertised in IGMP queries. Changes to this setting are valid only for IGMP Version 2.
- Group Limit—Displays the maximum number of groups permitted on an interface.
- Forward Interface—Displays the name of the interface that the selected interface forwards IGMP host reports to.

- Edit—Opens the Configure IGMP Parameters dialog box for the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Configure IGMP Parameters

Configuration > Routing > Multicast > IGMP > Protocol > Configure IGMP Parameters

The Configure IGMP Parameters dialog box lets you disable IGMP and change IGMP parameters on the selected interface.

Fields

- Interface—Displays the name of the interface being configured. You cannot change the information displayed in this field.
- Enable IGMP—Check this check box to enable IGMP on the interface. Uncheck the check box to disable IGMP on the interface. If you enabled multicast routing on the security appliance, then IGMP is enabled by default.
- Version—Choose the version of IGMP to enable on the interface. Choose 1 to enable IGMP Version 1, or 2 to enable IGMP Version 2. Some feature require IGMP Version 2. By default, the security appliance uses IGMP Version 2.
- Query Interval—Enter the interval, in seconds, at which the designated router sends IGMP host-query messages. Valid values range from 1 to 3600 seconds. The default value is 125 seconds.
- Query Timeout—Enter the period of time, in seconds, before which the security appliance takes over as the querier for the interface after the previous querier has stopped doing so. Valid values range from 60 to 300 seconds. The default value is 255 seconds.

- Response Time—Enter the maximum response time, in seconds, advertised in IGMP queries. If the security appliance does not receive any host reports within the designated response time, the IGMP group is pruned. Decreasing this value lets the security appliance prune groups faster. Valid values range from 1 to 12 seconds. The default value is 10 seconds. Changing this value is only valid only for IGMP Version 2.
- Group Limit—Enter the maximum number of host that can join on an interface. Valid values range from 1 to 500. The default value is 500.
- Forward Interface—Choose the name of an interface to forward IGMP host reports to. Choose "None" to disable host report forwarding. By default, host reports are not forwarded.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Static Group

Configuration > Routing > Multicast > IGMP > Static Group

Sometimes, hosts on a network may have a configuration that prevents them from answering IGMP queries. However, you still want multicast traffic to be forwarded to that network segment. There are two methods to pull multicast traffic down to a network segment:

- Use the Join Group pane to configure the interface as a member of the multicast group. With this method, the security appliance accepts the multicast packets in addition to forwarding them to the specified interface.
- Use the Static Group pane configure the security appliance to be a statically connected member of a group. With this method, the security appliance does not accept the packets itself, but only forwards them. Therefore, this method allows fast switching. The outgoing interface appears in the IGMP cache, but itself is not a member of the multicast group.

Fields

- Static Group—Displays the statically assigned multicast groups for each interface.
  - Interface—Displays the name of the security appliance interface.
  - Multicast Group Address—Displays the address of a multicast group assigned to the interface.
- Add—Opens the Add/Edit IGMP Static Group dialog box. Use this button to assign a new static group to an interface.
- Edit—Opens the Add/Edit IGMP Static Group dialog box. Use this button to edit an existing static group membership.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit IGMP Static Group

Configuration > Routing > Multicast > IGMP > Static Group > Add/Edit IGMP Static Group

Use the Add IGMP Static Group dialog box to statically assign a multicast group to an interface. Use the Edit IGMP Static Group dialog box to change existing static group assignments.

Fields

- Interface—Choose the name of the security appliance interface that you are configuring a multicast group for. If you are editing an existing entry, you cannot change this value.
- Multicast Group Address—Enter the address of a multicast group in this field. The group address must be from 224.0.0.0 to 239.255.255.255.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Multicast Route

Configuration > Routing > Multicast > MRoute

Defining static multicast routes lets you separate multicast traffic from unicast traffic. For example, when a path between a source and destination does not support multicast routing, the solution is to configure two multicast devices with a GRE tunnel between them and to send the multicast packets over the tunnel.

Static multicast routes are local to the security appliance and are not advertised or redistributed.

Fields

- Multicast Route—Displays the statically-defined multicast routes on the security appliance. Double-clicking an entry in the table opens the Add/Edit Multicast Route dialog box for that entry.
  - Source Address—Displays the IP address and mask, in CIDR notation, of the multicast source.
  - Source Interface—Displays the incoming interface for the multicast route.
  - Destination Interface—Displays the outgoing interface for the multicast route.
  - Admin Distance—Displays the administrative distance of the static multicast route.
- Add—Opens the Add/Edit Multicast Route dialog box. Use this button to add a new static route.

- Edit—Opens the Add/Edit Multicast Route dialog box. Use this button to change the selected static multicast route.
- Delete—Use this button to remove the selected static route.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Multicast Route

Configuration > Routing > Multicast > MRoute > Add/Edit Multicast Route

Use the Add Multicast Route dialog box to add a new static multicast route to the security appliance. Use the Edit Multicast Route dialog box to change an existing static multicast route.

Fields

- Source Address—Enter the IP address of the multicast source. You cannot change this value when editing an exiting static multicast route.
- Source Mask—Enter the network mask for the IP address of the multicast source or chose a common mask from the list. You cannot change this value when editing an exiting static multicast route.
- Source Interface—Choose the incoming interface for the multicast route.
- Destination Interface—(Optional) Choose the outgoing interface for the multicast route. If you specify the destination interface, the route is forwarded through the selected interface. If you do not choose a destination interface, then RPF is used to forward the route.
- Admin Distance—Enter the administrative distance of the static multicast route. If the static multicast route has the same administrative distance as the unicast route, then the static multicast route takes precedence.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

MBoundary

Configuration > Routing > Multicast > MBoundary

The MBoundary pane lets you configure a multicast boundary for administratively-scoped multicast addresses. A multicast boundary restricts multicast data packet flows and enables reuse of the same multicast group address in different administrative domains. When a multicast boundary is defined on an interface, only the multicast traffic permitted by the filter ACL passes through the interface.

Fields

The Multicast Boundary table contains the following information. Double-click a table entry to edit the multicast boundary filter settings.

- Interface—Lists the interfaces on the device.
- Boundary Filter—Lists the boundary filter entries for the specified interface. If a multicast boundary has not been defined for an interface, then this column displays "No Boundary Filters Configured" for the interface.
- AutoFilter—Shows if Auto-RP messages are denied by the boundary ACL. If the AutoFilter is enabled, the ACL also restricts the flow of Auto-RP messages. If the AutoFilter is disabled, all Auto-RP messages are passed by the interface. This setting is disabled by default.

You can perform the following actions on the entries of the Boundary table:

- Edit—Opens the Edit Boundary Filter dialog box.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Edit Boundary Filter

The Edit Boundary Filter dialog box displays the multicast boundary filter ACL. You can add and remove boundary filter ACL entries using this dialog box.

When the boundary filter configuration is applied to the security appliance, the ACL appears in the running configuration with the name *interface-name*_multicast, where the *interface-name* is the name of the interface the multicast boundary filter is applied to. If an ACL with that name already exists, a number is appended to the name, for example inside_multicast_1.

Fields

- Interface—Displays the interface for which you are configuring the multicast boundary filter ACL.
- Remove any Auto-RP group range—Check this check box to filter Auto-RP messages from sources denied by the boundary ACL. If not checked, all Auto-RP messages are passed.

The Boundary Filter table contains the following information:

- Action—The action for the filter entry. Permit allows the specified traffic to pass. Deny prevents the specified traffic from passing through the interface. When a multicast boundary filter is configured on an interface, multicast traffic is denied by default.
- Network Address—The multicast group address of the group being permitted or denied.

- Netmask—The network mask applied to the multicast group address.

You can perform the following actions on the Boundary Filter table:

- Insert—Inserts a neighbor filter entry before the selected entry.
- Add—Adds a neighbor filter entry after the selected entry.
- Edit—Edits the selected boundary filter.
- Delete—Removes the selected neighbor filter entry.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit/Insert Neighbor Filter Entry

The Add/Edit/Insert Neighbor Filter Entry dialog box lets you create the ACL entries for the multicast boundary ACL.

Fields

- Action—Select Permit or Deny for the neighbor filter ACL entry. Selecting Permit allows the multicast group advertisements through the interface. Selecting Deny prevents the specified multicast group advertisements from passing through the interface. When a multicast boundary is configured on an interface, all multicast traffic is prevented from passing through the interface unless permitted with a neighbor filter entry.
- Multicast Group Address—Enter the address of the multicast group being permitted or denied. Valid group addresses are from 224.0.0.0 to 239.255.255.255.
- Netmask—Type or select the netmask for the multicast group address.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

MForwarding

Configuration > Routing > Multicast >MForwarding

The MForwarding pane lets you disable and reenable multicast forwarding on a per interface basis. By default, multicast forwarding is enabled on all interfaces.

When multicast forwarding is disabled on an interface, the interface does not accept any multicast packets unless specifically configured through other methods. IGMP packets are also prevented when multicast forwarding is disabled.

Fields

- The Multicast Forwarding table displays the following information:
  - Interface—Displays the interfaces configured on the security appliance. Click an interface name to select the interface. Double-click an interface name to toggle the Multicast Forwarding Enabled status for the interface.
  - Multicast Forwarding Enabled—Displays Yes if multicast forwarding is enabled on the specified interface. Displays No if multicast forwarding is disabled on the specified interface. Double-click this entry to toggle Yes/No for the selected interface.
- Enable—Enables multicast forwarding on the selected interface.
- Disable—Disables multicast forwarding on the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

For More Information

- Configuring Multicast Routing, page 15-1

PIM

Routers use PIM to maintaining forwarding tables for forwarding multicast datagrams.

When you enable multicast routing on the security appliance, PIM is enabled on all interfaces by default. You can disable PIM on a per-interface basis.

For more information about configuring PIM, see the following:

- Protocol
- Neighbor Filter, page 15-13
- Bidirectional Neighbor Filter, page 15-14
- Rendezvous Points
- Route Tree
- Request Filter

Protocol

Configuration > Routing > Multicast > PIM > Protocol

The Protocol pane displays the interface-specific PIM properties.

Fields

- Protocol—Displays the PIM settings for each interface. Double-clicking an entry in the table opens the Edit PIM Protocol dialog box for that entry.
  - Interface—Displays the name of the security appliance interfaces.
  - PIM Enabled—Displays "Yes" if PIM is enabled on the interface, "No" if PIM is not enabled.
  - DR Priority—Displays the interface priority.
  - Hello Interval—Displays the frequency, in seconds, at which the interface sends PIM hello messages.
  - Join-Prune Interval—Displays the frequency, in seconds, at which the interface sends PIM join and prune advertisements.
- Edit—Opens the Edit PIM Protocol dialog box for the selected entry.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Edit PIM Protocol

Configuration > Routing > Multicast > PIM > Protocol > Edit PIM Protocol

The Edit PIM Protocol dialog box lets you change the PIM properties for the selected interface.

Fields

- Interface—*Display only.* Displays the name of the selected interface. You cannot edit this value.
- PIM Enabled—Check this check box to enable PIM on the selected interface. Uncheck this check box to disable PIM on the selected interface.
- DR Priority—Sets the designated router priority for the selected interface. The router with the highest DR priority on subnet becomes the designated router. Valid values range from 0 to 4294967294. The default DR priority is 1. Setting this value to 0 makes the security appliance interface ineligible to become the default router.
- Hello Interval—Enter the frequency, in seconds, at which the interface sends PIM hello messages. Valid values range from 1 to 3600 seconds. The default value is 30 seconds.
- Join-Prune Interval—Enter the frequency, in seconds, at which the interface sends PIM join and prune advertisements. Valid values range from 10 to 600 seconds. The default value is 60 seconds.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Neighbor Filter

The Neighbor Filter pane displays the PIM neighbor filters, if any, that are configured on the security appliance. A PIM neighbor filter is an ACL that defines the neighbor devices that can participate in PIM. If a neighbor filter is not configured for an interface, then there are no restrictions. If a PIM neighbor filter is configured, only those neighbors permitted by the filter list can participate in PIM with the security appliance.

When a PIM neighbor filter configuration is applied to the security appliance, an ACL appears in the running configuration with the name *interface-name*_multicast, where the *interface-name* is the name of the interface the multicast boundary filter is applied to. If an ACL with that name already exists, a number is appended to the name, for example inside_multicast_1. This ACL defines which devices can become PIM neighbors of the security appliance.

Fields

The PIM Neighbor Filter table displays the following information. Double-clicking an entry in the table opens the Edit Neighbor Filter Entry dialog box for the selected entry.

- Interface—Displays the name of the interface the PIM neighbor filter entry applies to.
- Action—Display “permit” if the specified neighbors are allowed to participate in PIM. Displays “deny” if the specified neighbors are prevented from participating in PIM.
- Network Address—The network address of the neighbor or neighbors being permitted or denied.
- Netmask—The network mask to use with the Network Address.

You can perform the following actions:

- Insert—Click to insert a neighbor filter entry before the selected entry.
- Add—Click to add a neighbor filter entry after the selected entry.
- Edit—Click to edit the selected neighbor filter entry.
- Delete—Click to remove the selected neighbor filter entry.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

For More Information

Add/Edit/Insert Neighbor Filter Entry, page 15-14

Add/Edit/Insert Neighbor Filter Entry

The Add/Edit/Insert Neighbor Filter Entry lets you create ACL entries for the PIM neighbor filter ACL.

Fields

- Interface—Select the name of the interface the PIM neighbor filter entry applies to from the list.
- Action—Select “permit” to allow the specified neighbors to participate in PIM. Select “deny” to prevent the specified neighbors from participating in PIM.
- Network Address—The network address of the neighbor or neighbors being permitted or denied.
- Netmask—The network mask to use with the Network Address.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Bidirectional Neighbor Filter

The Bidirectional Neighbor Filter pane shows the PIM bidirectional neighbor filters, if any, that are configured on the security appliance. A PIM bidirectional neighbor filters is an ACL that defines the neighbor devices that can participate in the DF election. If a PIM bidirectional neighbor filter is not configured for an interface, then there are no restrictions. If a PIM bidirectional neighbor filter is configured, only those neighbors permitted by the ACL can participate in DF election process.

When a PIM bidirectional neighbor filter configuration is applied to the security appliance, an ACL appears in the running configuration with the name *interface-name*_multicast, where the *interface-name* is the name of the interface the multicast boundary filter is applied to. If an ACL with that name already exists, a number is appended to the name, for example inside_multicast_1. This ACL defines which devices can become PIM neighbors of the security appliance.

Bidirectional PIM allows multicast routers to keep reduced state information. All of the multicast routers in a segment must be bidirectionally enabled for bidir to elect a DF.

The PIM bidirectional neighbor filters enable the transition from a sparse-mode-only network to a bidir network by letting you specify the routers that should participate in DF election while still allowing all routers to participate in the sparse-mode domain. The bidir-enabled routers can elect a DF from among themselves, even when there are non-bidir routers on the segment. Multicast boundaries on the non-bidir routers prevent PIM messages and data from the bidir groups from leaking in or out of the bidir subset cloud.

When a PIM bidirectional neighbor filter is enabled, the routers that are permitted by the ACL are considered to be bidir-capable. Therefore:

- If a permitted neighbor does not support bidir, the DF election does not occur.

- If a denied neighbor supports bidir, then DF election does not occur.
- If a denied neighbor does not support bidir, the DF election can occur.

Fields

The PIM Bidirectional Neighbor Filter table contains the following entries. Double-click an entry to open the Edit Bidirectional Neighbor Filter Entry dialog box for that entry.

- Interface—Displays the interface the bidirectional neighbor filter applies to.
- Action—Displays "permit" if the bidirectional neighbor filter entry allows participation in the DF election process. Display "deny" if the entry prevents the specified addresses from participating in the DF election process.
- Network Address—The address being permitted or denied.
- Netmask—The network mask to apply to the Network Address.

You can perform the following actions:

- Insert—Click to insert a bidirectional neighbor filter entry before the selected entry.
- Add—Click to add a bidirectional neighbor filter entry after the selected entry.
- Edit—Click to edit the selected bidirectional neighbor filter entry.
- Delete—Click to remove the selected bidirectional neighbor filter entry.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

For More Information

Add/Edit/Insert Bidirectional Neighbor Filter Entry, page 15-15

Add/Edit/Insert Bidirectional Neighbor Filter Entry

The Add/Edit/Insert Bidirectional Neighbor Filter Entry dialog box lets you create ACL entries for the PIM bidirectional neighbor filter ACL.

Fields

- Interface—Select the interface for which you are configuring the PIM bidirectional neighbor filter ACL entry.
- Action—Select permit to allow the specified devices to participate in the DF election. Select deny to prevent the specified devices from participating in the DF election.
- Network Address—The network address of the neighbor or neighbors being permitted or denied.
- Netmask—The network mask to use with the Network Address.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Rendezvous Points

Configuration > Routing > Multicast > PIM > Rendezvous Points

When you configure PIM, you must choose one or more routers to operate as the RP. An RP is a single, common root of a shared distribution tree and is statically configured on each router. First hop routers use the RP to send register packets on behalf of the source multicast hosts.

You can configure a single RP to serve more than one group. If a specific group is not specified, the RP for the group is applied to the entire IP multicast group range (224.0.0.0/4).

You can configure more than one RP, but you cannot have more than one entry with the same RP.

Fields

- Generate IOS compatible register messages—Check this check box if your RP is a Cisco IOS router. The security appliance software accepts register messages with the checksum on the PIM header and only the next 4 bytes rather than using the Cisco IOS software method—accepting register messages with the checksum on the entire PIM message for all PIM message types.
- Rendezvous Points—Displays the RPs configured on the security appliance.
  - Rendezvous Point—Displays the IP address of the RP.
  - Multicast Groups—Displays the multicast groups associated with the RP. Displays "--All Groups--" if the RP is associated with all multicast groups on the interface.
  - Bi-directional—Displays "Yes" if the specified multicast groups are to operate in bidirectional mode. Displays "No" if the specified groups are to operate in sparse mode.
- Add—Opens the Add/Edit Rendezvous Point dialog box. Use this button to add a new RP entry.
- Edit—Opens the Add/Edit Rendezvous Point dialog box. Use this button to change an existing RP entry.
- Delete—Removes the selected RP entry from the Rendezvous Point table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Rendezvous Point

Configuration > Routing > Multicast > PIM > Rendezvous Points >Add/Edit Rendezvous Point

The Add Rendezvous Point dialog box lets you add a new entry to the Rendezvous Point table. The Edit Rendezvous Point dialog box lets you change an existing RP entry.

Restrictions

- You cannot use the same RP address twice.
- You cannot specify All Groups for more than one RP.

Fields

- Rendezvous Point IP Address—Enter the IP address of the RP. This is a unicast address. When editing an existing RP entry, you cannot change this value.
- Use bi-directional forwarding—Check this check box if you want the specified multicast groups to operation in bidirectional mode. In bidirectional mode, if the security appliance receives a multicast packet and has no directly connected members or PIM neighbors present, it sends a Prune message back to the source. Uncheck this check box if you want the specified multicast groups to operate in sparse mode.

Note The security appliance always advertises the bidir capability in the PIM hello messages regardless of the actual bidir configuration.

- Use this RP for All Multicast Groups—Choose this option to use the specified RP for all multicast groups on the interface.
- Use this RP for the Multicast Groups as specified below—Choose this option to designate the multicast groups to use with specified RP.
- Multicast Groups—Displays the multicast groups associated with the specified RP.

  The table entries are processed from the top down. You can create an RP entry that includes a range of multicast groups but excludes specific groups within that range by placing deny rules for the specific groups at the top of the table and the permit rule for the range of multicast groups below the deny statements.

  Double-click an entry to open the Multicast Group dialog box for the selected entry.

  - Action—Displays “Permit” if the multicast group is included or “deny” if the multicast group is excluded.
  - Multicast Group Address—Displays the address of the multicast group.
  - Netmask—Displays the network mask of the multicast group address.
- Insert Before—Opens the Multicast Group dialog box. Use this button to add a new multicast group entry before the selected entry in the table.
- Insert After—Opens the Multicast Group dialog box. Use this button to add a new multicast group entry after the selected entry in the table.
- Add—Opens the Multicast Group dialog box. Use this button to add a new multicast group entry at the bottom of the table.
- Edit—Opens the Multicast Group dialog box. Use this button to change the information for the selected multicast group entry.
- Delete—Removes the selected multicast group entry from the table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Multicast Group

Configuration > Routing > Multicast > PIM > Rendezvous Points > Add/Edit Rendezvous Point > Multicast Group (You can get to this dialog through various paths.)

Multicast groups are lists of access rules that define which multicast addresses are part of the group. A multicast group can contain a single multicast address or a range of multicast addresses. Use the Add Multicast Group dialog box to create a new multicast group rule. Use the Edit Multicast Group dialog box to modify an existing multicast group rule.

Fields

- Action—Choose “Permit” to create a group rule that allows the specified multicast addresses; choose “Deny” to create a group rule that filters the specified multicast addresses.
- Multicast Group Address—Enter the multicast address associated with the group.
- Netmask—Enter or choose the network mask for the multicast group address.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Request Filter

Configuration > Routing > Multicast > PIM > Request Filter

When the security appliance is acting as an RP, you can restrict specific multicast sources from registering with it. This prevents unauthorized sources from registering with the RP. The Request Filter pane lets you define the multicast sources from which the security appliance will accept PIM register messages.

Fields

- Multicast Groups—Displays the request filter access rules.

The table entries are processed from the top down. You can create an entry that includes a range of multicast groups but excludes specific groups within that range by placing deny rules for the specific groups at the top of the table and the permit rule for the range of multicast groups below the deny statements.

Double-click an entry to open the Request Filter Entry dialog box for the selected entry.

- Action—Displays "Permit" if the multicast source is allowed to register or "deny" if the multicast source is excluded.
- Source—Displays the address of the source of the register message.
- Destination—Displays the multicast destination address.

- Insert Before—Opens the Request Filter Entry dialog box. Use this button to add a new multicast group entry before the selected entry in the table.
- Insert After—Opens the Request Filter Entry dialog box. Use this button to add a new multicast group entry after the selected entry in the table.
- Add—Opens the Request Filter Entry dialog box. Use this button to add a new multicast group entry at the bottom of the table.
- Edit—Opens the Request Filter Entry dialog box. Use this button to change the information for the selected multicast group entry.
- Delete—Removes the selected multicast group entry from the table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Request Filter Entry

Configuration > Routing > Multicast > PIM > Request Filter > Request Filter Entry

The Request Filter Entry dialog box lets you define the multicast sources that are allowed to register with the security appliance when the security appliance acts as an RP. You create the filter rules based on the source IP address and the destination multicast address.

Fields

- Action—Choose "Permit" to create a rule that allows the specified source of the specified multicast traffic to register with the security appliance; choose "Deny" to create a rule that prevents the specified source of the specified multicast traffic from registering with the security appliance.
- Source IP Address—Enter the IP address for the source of the register message.
- Source Netmask—Enter or choose the network mask for the source of the register message.
- Destination IP Address—Enter the multicast destination address.
- Destination Netmask—Enter or choose the network mask for the multicast destination address.

Modes

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Route Tree

Configuration > Routing > Multicast > PIM > Route Tree

By default, PIM leaf routers join the shortest-path tree immediately after the first packet arrives from a new source. This reduces delay, but requires more memory than shared tree.

You can configure whether the security appliance should join shortest-path tree or use shared tree, either for all multicast groups or only for specific multicast addresses.

Fields

- Use Shortest Path Tree for All Groups—Choose this option to use shortest-path tree for all multicast groups.
- Use Shared Tree for All Groups—Choose this option to use shared tree for all multicast groups.
- Use Shared Tree for the Groups specified below—Choose this option to use shared tree for the groups specified in the Multicast Groups table. Shortest-path tree is used for any group not specified in the Multicast Groups table.
- Multicast Groups—Displays the multicast groups to use Shared Tree with.

  The table entries are processed from the top down. You can create an entry that includes a range of multicast groups but excludes specific groups within that range by placing deny rules for the specific groups at the top of the table and the permit rule for the range of multicast groups below the deny statements.

  Double-click an entry to open the Multicast Group dialog box for the selected entry.

  - Action—Displays “Permit” if the multicast group is included or “deny” if the multicast group is excluded.
  - Multicast Group Address—Displays the address of the multicast group.
  - Netmask—Displays the network mask of the multicast group address.
- Insert Before—Opens the Multicast Group dialog box. Use this button to add a new multicast group entry before the selected entry in the table.
- Insert After—Opens the Multicast Group dialog box. Use this button to add a new multicast group entry after the selected entry in the table.
- Add—Opens the Multicast Group dialog box. Use this button to add a new multicast group entry at the bottom of the table.
- Edit—Opens the Multicast Group dialog box. Use this button to change the information for the selected multicast group entry.
- Delete—Removes the selected multicast group entry from the table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

CHAPTER 16

Firewall Mode Overview

This chapter describes how the firewall works in each firewall mode. To set the mode at the CLI, see the "Setting Transparent or Routed Firewall Mode at the CLI" section on page 2-5.

The security appliance can run in two firewall modes:

- Routed mode
- Transparent mode

In routed mode, the security appliance is considered to be a router hop in the network. It can perform NAT between connected networks, and can use OSPF or passive RIP (in single context mode). Routed mode supports many interfaces. Each interface is on a different subnet. You can share interfaces between contexts.

In transparent mode, the security appliance acts like a "bump in the wire," or a "stealth firewall," and is not a router hop. The security appliance connects the same network on its inside and outside interfaces. No dynamic routing protocols or NAT are used. However, like routed mode, transparent mode also requires access lists to allow any traffic through the security appliance, except for ARP packets, which are allowed automatically. Transparent mode can allow certain types of traffic in an access list that are blocked by routed mode, including unsupported routing protocols. Transparent mode can also optionally use EtherType access lists to allow non-IP traffic. Transparent mode only supports two interfaces, an inside interface and an outside interface, in addition to a dedicated management interface, if available for your platform.

Note In transparent firewall mode, you do not set the IP address for each interface, but rather for the whole security appliance or context. (The exception is for the Management 0/0 management-only interface, for which you can set an IP address; this interface does not pass through traffic.) The security appliance uses the overall management IP address as the source address for packets originating on the security appliance. The management IP address must be on the same subnet as the connected network.

This chapter includes the following sections:

Routed Mode Overview

- How Data Moves Through the Security Appliance in Routed Firewall Mode, page 16-3

IP Routing Support

The security appliance acts as a router between connected networks, and each interface requires an IP address on a different subnet. In single context mode, the routed firewall supports OSPF and RIP (in passive mode). Multiple context mode supports static routes only. We recommend using the advanced routing capabilities of the upstream and downstream routers instead of relying on the security appliance for extensive routing needs.

Network Address Translation

NAT substitutes the local address on a packet with a global address that is routable on the destination network. By default, NAT is not required. If you want to enforce a NAT policy that requires hosts on a higher security interface (inside) to use NAT when communicating with a lower security interface (outside), you can enable NAT control (see the **nat-control** command).

Note NAT control was the default behavior for software versions earlier than Version 7.0. If you upgrade a security appliance from an earlier version, then the **nat-control** command is automatically added to your configuration to maintain the expected behavior.

Some of the benefits of NAT include the following:

- You can use private addresses on your inside networks. Private addresses are not routable on the Internet.
- NAT hides the local addresses from other networks, so attackers cannot learn the real address of a host.
- NAT can resolve IP routing problems by supporting overlapping IP addresses.

Figure 16-1 shows a typical NAT scenario, with a private network on the inside. When the inside user sends a packet to a web server on the Internet, the local source address of the packet is changed to a routable global address. When the web server responds, it sends the response to the global address, and the security appliance receives the packet. The security appliance then translates the global address to the local address before sending it on to the user.

Figure 16-1 NAT Example

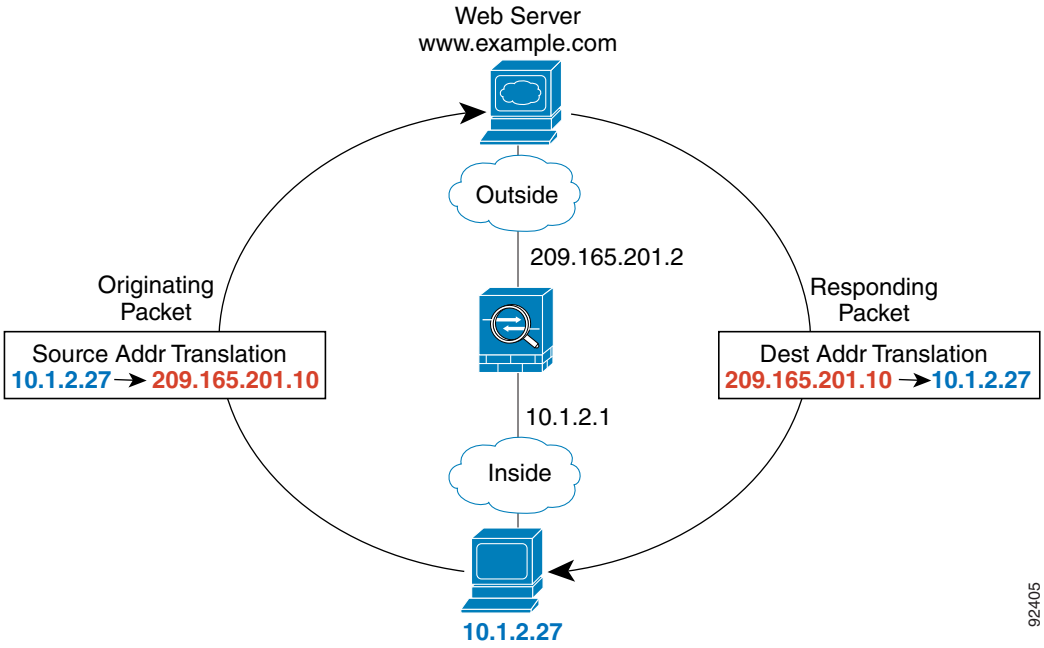

How Data Moves Through the Security Appliance in Routed Firewall Mode

This section describes how data moves through the security appliance in routed firewall mode, and includes the following topics:

An Inside User Visits a Web Server

Figure 16-2 shows an inside user accessing an outside web server.

Figure 16-2 Inside to Outside

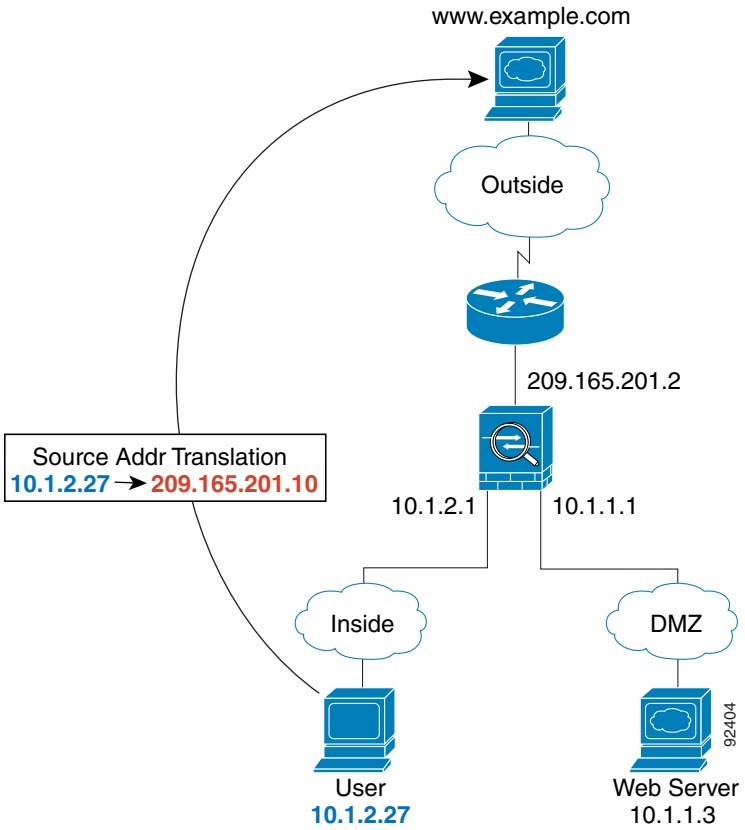

The following steps describe how data moves through the security appliance (see Figure 16-2):

1. The user on the inside network requests a web page from www.example.com.
2. The security appliance receives the packet and because it is a new session, the security appliance verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).

   For multiple context mode, the security appliance first classifies the packet according to either a unique interface or a unique destination address associated with a context; the destination address is associated by matching an address translation in a context. In this case, the interface would be unique; the www.example.com IP address does not have a current address translation in a context.
3. The security appliance translates the local source address (10.1.2.27) to the global address 209.165.201.10, which is on the outside interface subnet.

   The global address could be on any subnet, but routing is simplified when it is on the outside interface subnet.
4. The security appliance then records that a session is established and forwards the packet from the outside interface.

5. When www.example.com responds to the request, the packet goes through the security appliance, and because the session is already established, the packet bypasses the many lookups associated with a new connection. The security appliance performs NAT by translating the global destination address to the local user address, 10.1.2.27.
6. The security appliance forwards the packet to the inside user.

An Outside User Visits a Web Server on the DMZ

Figure 16-3 shows an outside user accessing the DMZ web server.

Figure 16-3 Outside to DMZ

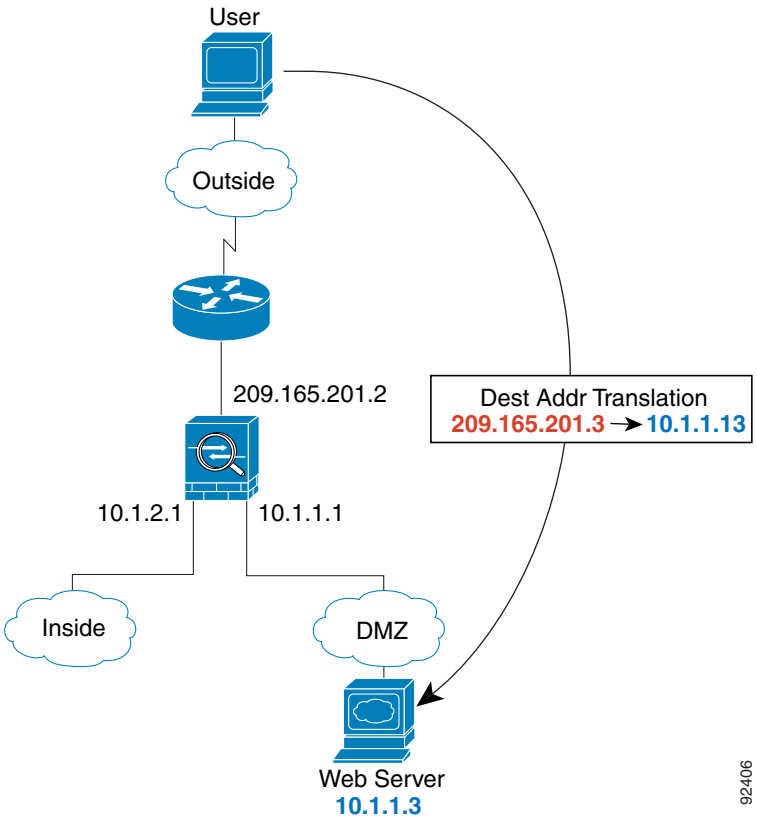

The following steps describe how data moves through the security appliance (see Figure 16-3):

1. A user on the outside network requests a web page from the DMZ web server using the global destination address of 209.165.201.3, which is on the outside interface subnet.
2. The security appliance receives the packet and because it is a new session, the security appliance verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).

   For multiple context mode, the security appliance first classifies the packet according to either a unique interface or a unique destination address associated with a context; the destination address is associated by matching an address translation in a context. In this case, the classifier "knows" that the DMZ web server address belongs to a certain context because of the server address translation.
3. The security appliance translates the destination address to the local address 10.1.1.3.

4. The security appliance then adds a session entry to the fast path and forwards the packet from the DMZ interface.
5. When the DMZ web server responds to the request, the packet goes through the security appliance and because the session is already established, the packet bypasses the many lookups associated with a new connection. The security appliance performs NAT by translating the local source address to 209.165.201.3.
6. The security appliance forwards the packet to the outside user.

An Inside User Visits a Web Server on the DMZ

Figure 16-4 shows an inside user accessing the DMZ web server.

Figure 16-4 Inside to DMZ

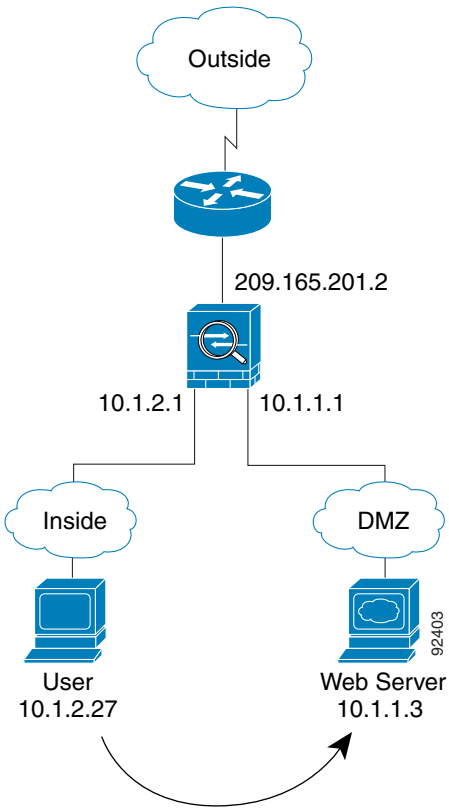

The following steps describe how data moves through the security appliance (see Figure 16-4):

1. A user on the inside network requests a web page from the DMZ web server using the destination address of 10.1.1.3.
2. The security appliance receives the packet and because it is a new session, the security appliance verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).

   For multiple context mode, the security appliance first classifies the packet according to either a unique interface or a unique destination address associated with a context; the destination address is associated by matching an address translation in a context. In this case, the interface is unique; the web server IP address does not have a current address translation.

3. The security appliance then records that a session is established and forwards the packet out of the DMZ interface.
4. When the DMZ web server responds to the request, the packet goes through the fast path, which lets the packet bypass the many lookups associated with a new connection.
5. The security appliance forwards the packet to the inside user.

An Outside User Attempts to Access an Inside Host

Figure 16-5 shows an outside user attempting to access the inside network.

Figure 16-5 Outside to Inside

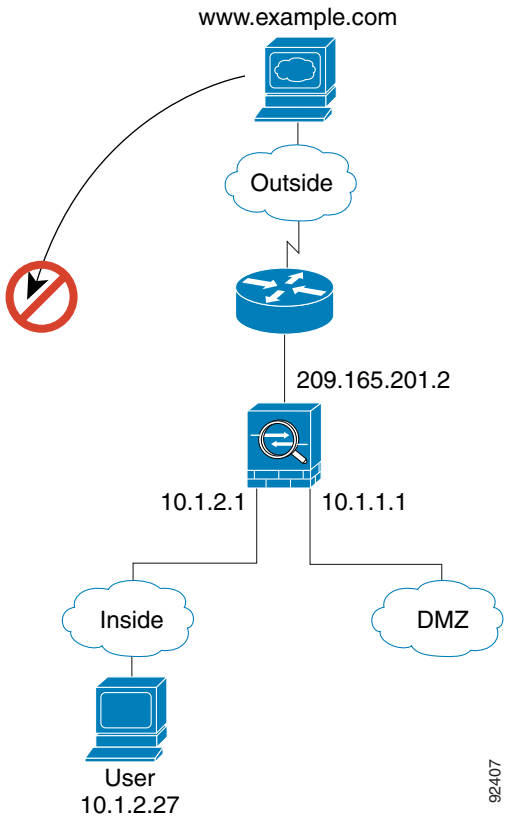

The following steps describe how data moves through the security appliance (see Figure 16-5):

1. A user on the outside network attempts to reach an inside host (assuming the host has a routable IP address).

   If the inside network uses private addresses, no outside user can reach the inside network without NAT. The outside user might attempt to reach an inside user by using an existing NAT session.
2. The security appliance receives the packet and because it is a new session, the security appliance verifies if the packet is allowed according to the security policy (access lists, filters, AAA).
3. The packet is denied, and the security appliance drops the packet and logs the connection attempt.

   If the outside user is attempting to attack the inside network, the security appliance employs many technologies to determine if a packet is valid for an already established session.

A DMZ User Attempts to Access an Inside Host

Figure 16-6 shows a user in the DMZ attempting to access the inside network.

Figure 16-6 DMZ to Inside

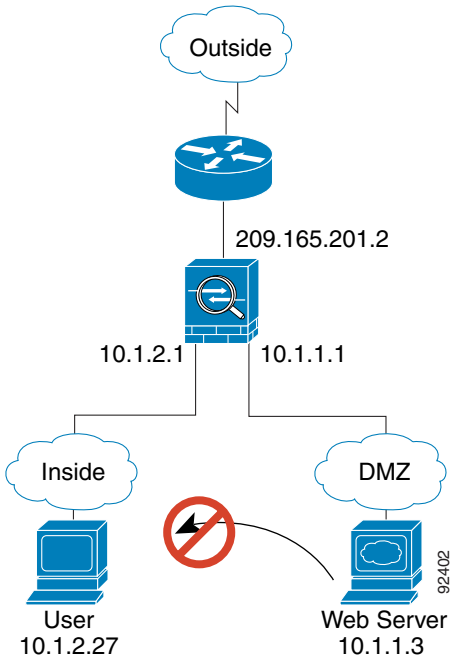

The following steps describe how data moves through the security appliance (see Figure 16-6):

1. A user on the DMZ network attempts to reach an inside host. Because the DMZ does not have to route the traffic on the internet, the private addressing scheme does not prevent routing.
2. The security appliance receives the packet and because it is a new session, the security appliance verifies if the packet is allowed according to the security policy (access lists, filters, AAA).
3. The packet is denied, and the security appliance drops the packet and logs the connection attempt.

Transparent Mode Overview

This section describes transparent firewall mode, and includes the following topics:

Transparent Firewall Features

Traditionally, a firewall is a routed hop and acts as a default gateway for hosts that connect to one of its screened subnets. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a "bump in the wire," or a "stealth firewall," and is not seen as a router hop to connected devices. The security appliance connects the same network on its inside and outside ports. Because the firewall is not a routed hop, you can easily introduce a transparent firewall into an existing network; IP readdressing is unnecessary.

Maintenance is facilitated because there are no complicated routing patterns to troubleshoot and no NAT configuration.

Even though transparent mode acts as a bridge, Layer 3 traffic, such as IP traffic, cannot pass through the security appliance unless you explicitly permit it with an extended access list. The only traffic allowed through the transparent firewall without an access list is ARP traffic. ARP traffic can be controlled by ARP inspection.

In routed mode, some types of traffic cannot pass through the security appliance even if you allow it in an access list. The transparent firewall, however, can allow almost all traffic through using either an extended access list (for IP traffic) or an EtherType access list (for non-IP traffic). The following destination MAC addresses are allowed through the transparent firewall. Any MAC address not on this list is dropped.

- TRUE broadcast destination MAC address equal to FFFF.FFFF.FFFF
- IPv4 multicast MAC addresses from 0100.5E00.0000 to 0100.5EFE.FFFF
- IPv6 multicast MAC addresses from 3333.0000.0000 to 3333.FFFF.FFFF
- BPDU multicast address equal to 0100.0CCC.CCCD
- Appletalk multicast MAC addresses from 0900.0700.0000 to 0900.07FF.FFFF

Note The transparent mode security appliance does not pass CDP packets.

For example, you can establish routing protocol adjacencies through a transparent firewall; you can allow OSPF, RIP, EIGRP, or BGP traffic through based on an extended access list. Likewise, protocols like HSRP or VRRP can pass through the security appliance.

Non-IP traffic (for example AppleTalk, IPX, BPDUs, and MPLS) can be configured to go through using an EtherType access list.

For features that are not directly supported on the transparent firewall, you can allow traffic to pass through so that upstream and downstream routers can support the functionality. For example, by using an extended access list, you can allow DHCP traffic (instead of the unsupported DHCP relay feature) or multicast traffic such as that created by IP/TV.

When the security appliance runs in transparent mode, the outgoing interface of a packet is determined by performing a MAC address lookup instead of a route lookup. Route statements can still be configured, but they only apply to security appliance-originated traffic. For example, if your syslog server is located on a remote network, you must use a static route so the security appliance can reach that subnet.

Using the Transparent Firewall in Your Network

Figure 16-7 shows a typical transparent firewall network where the outside devices are on the same subnet as the inside devices. The inside router and hosts appear to be directly connected to the outside router.

Figure 16-7 Transparent Firewall Network

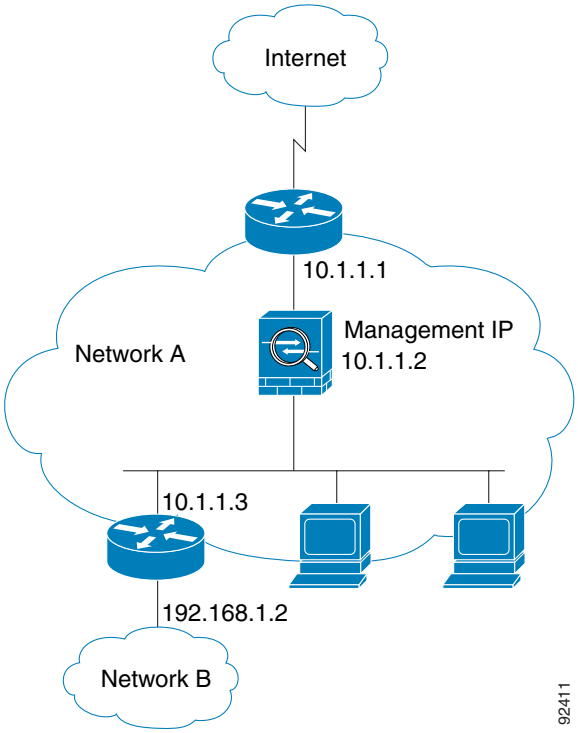

Transparent Firewall Guidelines

Follow these guidelines when planning your transparent firewall network:

- A management IP address is required; for multiple context mode, an IP address is required for each context.

  Unlike routed mode, which requires an IP address for each interface, a transparent firewall has an IP address assigned to the entire device. The security appliance uses this IP address as the source address for packets originating on the security appliance, such as system messages or AAA communications.

  The management IP address must be on the same subnet as the connected network. You cannot set the subnet to a host subnet (255.255.255.255).

  You can configure an IP address for the Management 0/0 management-only interface. This IP address can be on a separate subnet from the main management IP address.

- The transparent security appliance uses an inside interface and an outside interface only. If your platform includes a dedicated management interface, you can also configure the management interface or subinterface for management traffic only.

In single mode, you can only use two data interfaces (and the dedicated management interface, if available) even if your security appliance includes more than two interfaces.

- Each directly connected network must be on the same subnet.
- Do not specify the security appliance management IP address as the default gateway for connected devices; devices need to specify the router on the other side of the security appliance as the default gateway.
- For multiple context mode, each context must use different interfaces; you cannot share an interface across contexts.
- For multiple context mode, each context typically uses a different subnet. You can use overlapping subnets, but your network topology requires router and NAT configuration to make it possible from a routing standpoint.
- You must use an extended access list to allow Layer 3 traffic, such as IP traffic, through the security appliance.

  You can also optionally use an EtherType access list to allow non-IP traffic through.

Unsupported Features in Transparent Mode

Table 16-1 lists the features are not supported in transparent mode.

Table 16-1 Unsupported Features in Transparent Mode

Feature	Description
Dynamic DNS	—
DHCP relay	The transparent firewall can act as a DHCP server, but it does not support the DHCP relay commands. DHCP relay is not required because you can allow DHCP traffic to pass through using two extended access lists: one that allows DCHP requests from the inside interface to the outside, and one that allows the replies from the server in the other direction.
Dynamic routing protocols	You can, however, add static routes for traffic originating on the security appliance. You can also allow dynamic routing protocols through the security appliance using an extended access list.
IPv6	You also cannot allow IPv6 using an EtherType access list.
Multicast	You can allow multicast traffic through the security appliance by allowing it in an extended access list.
NAT	NAT is performed on the upstream router.
QoS	—
VPN termination for through traffic	The transparent firewall supports site-to-site VPN tunnels for management connections only. It does not terminate VPN connections for traffic through the security appliance. You can pass VPN traffic through the security appliance using an extended access list, but it does not terminate non-management connections. WebVPN is also not supported.

How Data Moves Through the Transparent Firewall

Figure 16-8 shows a typical transparent firewall implementation with an inside network that contains a public web server. The security appliance has an access list so that the inside users can access Internet resources. Another access list lets the outside users access only the web server on the inside network.

Figure 16-8 Typical Transparent Firewall Data Path

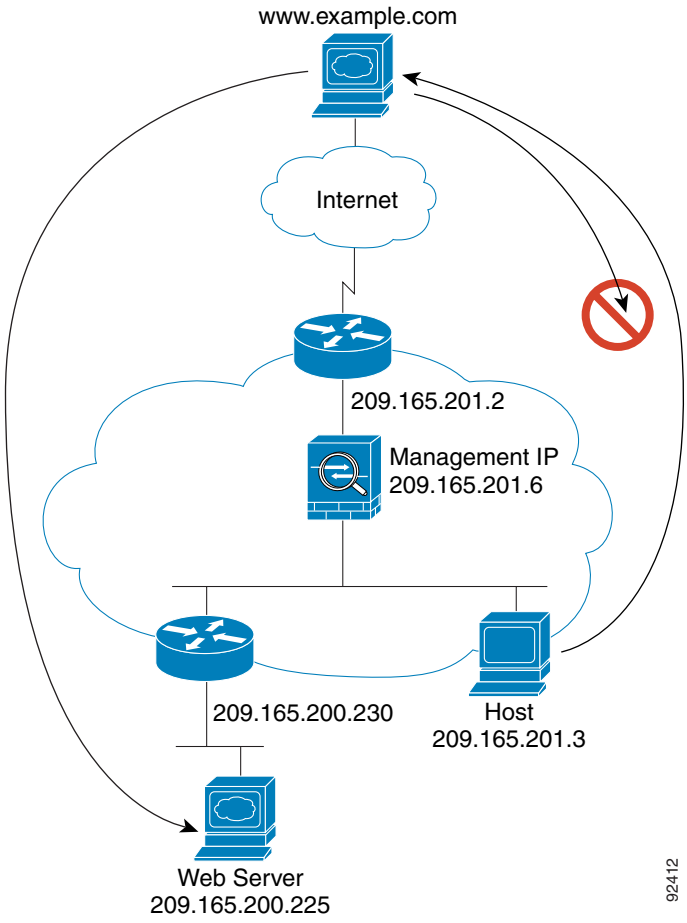

This section describes how data moves through the security appliance, and includes the following topics:

- An Inside User Visits a Web Server, page 16-13
- An Outside User Visits a Web Server on the Inside Network, page 16-14
- An Outside User Attempts to Access an Inside Host, page 16-15

An Inside User Visits a Web Server

Figure 16-9 shows an inside user accessing an outside web server.

Figure 16-9 Inside to Outside

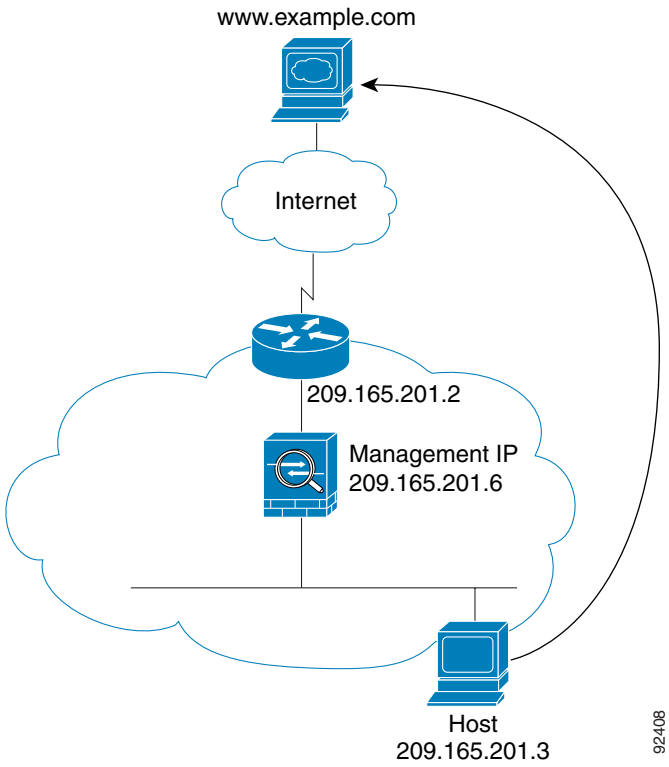

The following steps describe how data moves through the security appliance (see Figure 16-9):

1. The user on the inside network requests a web page from www.example.com.
2. The security appliance receives the packet and adds the source MAC address to the MAC address table, if required. Because it is a new session, it verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).

   For multiple context mode, the security appliance first classifies the packet according to a unique interface.
3. The security appliance and records that a session is established.
4. If the destination MAC address is in its table, the security appliance forwards the packet out of the outside interface. The destination MAC address is that of the upstream router, 209.186.201.2.

   If the destination MAC address is not in the security appliance table, the security appliance attempts to discover the MAC address by sending an ARP request and a ping. The first packet is dropped.
5. When the web server responds to the request, the security appliance adds the web server MAC address to the MAC address table, if required, and because the session is already established, the packet bypasses the many lookups associated with a new connection.
6. The security appliance forwards the packet to the inside user.

An Outside User Visits a Web Server on the Inside Network

Figure 16-10 shows an outside user accessing the inside web server.

Figure 16-10 Outside to Inside

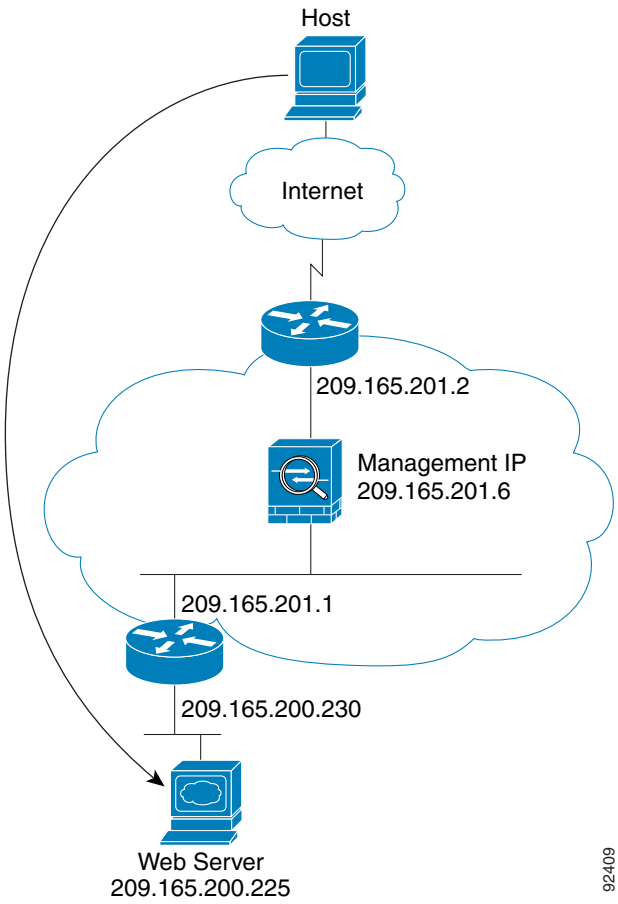

The following steps describe how data moves through the security appliance (see Figure 16-10):

1. A user on the outside network requests a web page from the inside web server.
2. The security appliance receives the packet and adds the source MAC address to the MAC address table, if required. Because it is a new session, it verifies that the packet is allowed according to the terms of the security policy (access lists, filters, AAA).

   For multiple context mode, the security appliance first classifies the packet according to a unique interface.
3. The security appliance records that a session is established.
4. If the destination MAC address is in its table, the security appliance forwards the packet out of the inside interface. The destination MAC address is that of the downstream router, 209.186.201.1.

   If the destination MAC address is not in the security appliance table, the security appliance attempts to discover the MAC address by sending an ARP request and a ping. The first packet is dropped.

5. When the web server responds to the request, the security appliance adds the web server MAC address to the MAC address table, if required, and because the session is already established, the packet bypasses the many lookups associated with a new connection.
6. The security appliance forwards the packet to the outside user.

An Outside User Attempts to Access an Inside Host

Figure 16-11 shows an outside user attempting to access a host on the inside network.

Figure 16-11 Outside to Inside

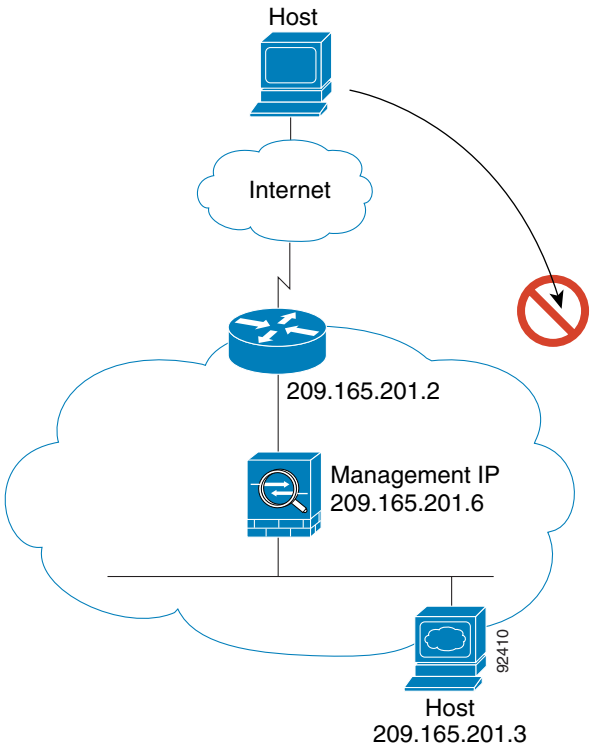

The following steps describe how data moves through the security appliance (see Figure 16-11):

1. A user on the outside network attempts to reach an inside host.
2. The security appliance receives the packet and adds the source MAC address to the MAC address table, if required. Because it is a new session, it verifies if the packet is allowed according to the terms of the security policy (access lists, filters, AAA).

   For multiple context mode, the security appliance first classifies the packet according to a unique interface.
3. The packet is denied, and the security appliance drops the packet.
4. If the outside user is attempting to attack the inside network, the security appliance employs many technologies to determine if a packet is valid for an already established session.

Configuring Access Rules

Access Rules

Configuration > Security Policy > Access Rules

The Access Rules window shows your entire network security policy expressed in rules.

When you choose the **Access Rules** option, this window lets you define access control lists to control the access of a specific host or network to another host/network, including the protocol or port that can be used. Conduits and outbound lists have been superceded by access lists.

By default on the security appliance, traffic from a higher security level (for example, inside) can access a lower security level (for example, outside): there is an implicit access list on the inside interface allowing all outbound IP traffic from the inside network. (The security appliance denies traffic destined for the inside network from the outside network using the Adaptive Security Algorithm. Adaptive Security Algorithm is a stateful approach to security. Every inbound packet is checked against the Adaptive Security Algorithm and against connection state information in memory.) The implicit access list appears in ASDM, but you cannot edit it. To limit outbound traffic, you can add an access list (in which case, the implicit access list is removed).

Every inbound packet is checked using the Adaptive Security Algorithm unless a connection is already established. By default on the security appliance, no traffic can pass through the firewall unless you add an access list to allow it.

To allow traffic that is normally denied by the Adaptive Security Algorithm, you can add an access list; for example, you can allow public access to a web server on a DMZ network by adding an access list to the outside interface.

Restrictions

At the end of each access list, there is an implicit, unwritten rule that denies all traffic that is not permitted. If traffic is not explicitly permitted by an access control entry, it will be denied. ACEs are referred to as rules in this topic.

Prerequisites

If desired, create network groups on the Addresses tab.

Fields

Note: You can adjust the table column widths by moving your cursor over a column line until it turns into a double arrow. Click and drag the column line to the desired size.

- Add—Adds a new access rule.
- Edit—Edits an access rule.

- Delete—Deletes an access rule.
- Move Up—Moves a rule up. Rules are assessed in the order they appear in this table, so the order can matter if you have overlapping rules.
- Move Down—Moves a rule down.
- Cut—Cuts a rule.
- Copy—Copies the parameters of a rule so you can start a new rule with the same parameters using the Paste button.
- Paste—Opens an Add/Edit Rule dialog box with the copied or cut parameters of a rule prefilled. You can then make any modifications and add it to the table. The Paste button adds the rule above the selected rule. The Paste After item, available from the Paste drop-down list, adds the rule after the selected rule.
- Find—Filters the display to show only matching rules. Clicking **Find** opens the Filter field. Click **Find** again to hide the Filter field.
  - Filter drop-down list—Choose the criteria to filter on, either Interface, Source, Destination, Service, or Rule Query. A rule query is a collection of multiple criteria that you can save and use repeatedly.
  - Filter field—For the Interface type, this field becomes a drop-down list so you can choose an interface name. For the Rule Query type, the drop-down list includes all defined rule queries. The Source and Destination types accept an IP address. You can type one manually, or browse for one by clicking the **...** button and launching the Browse Address dialog box. The Service type accepts a TCP, UDP, TCP-UDP, ICMP, or IP protocol type. You can type one manually, or browse for one by clicking the **...** button and launching the Browse Service Groups dialog box. The Filter field accepts multiple entries separated by a comma or space. Wildcards are also allowed.
  - Filter—Runs the filter.
  - Clear—Clears the matches and displays all.
  - Rule Query—Opens the Rule Queries dialog box so you can manage named rule queries.
- Show Rule Flow Diagram—Shows the Rule Flow Diagram area under the rule table. This diagram shows the networks, type of traffic, interface name, direction of flow, and action.
- Packet Trace—Opens the Packet Tracer tool with the parameters pre-filled with the characteristics of the selected rule.

The following description summarizes the columns in the Access Rules table. You can edit the contents of these columns by double-clicking on a table row. Rules are displayed in the order of execution. If you right-click a rule, you see all of the options represented by the buttons above, as well as Insert and Insert After items. These items either insert a new rule before the selected rule (Insert) or after the selected rule (Insert After.)

- No—Indicates the order of evaluation for the rule.
- Enabled—Indicates whether the rule is enabled or disabled.
- Source—Specifies the IP address, network object group, interface IP, or any, from which traffic is permitted or denied to the destination specified in the Destination Type field. An address column might contain an interface name with the word any, such as inside:any. This means that any host on the inside interface is affected by the rule.
- Destination—Specifies the IP address, network object group, interface IP, or any, to which traffic is permitted or denied from the source specified in the Source Type field. An address column might contain an interface name with the word any, such as outside:any. This means that any host on the

outside interface is affected by the rule. Also in detail mode, an address column might contain IP addresses in square brackets, for example [209.165.201.1-209.165.201.30]. These addresses are translated addresses. When an inside host makes a connection to an outside host, the firewall maps the address of the inside host to an address from the pool. After a host creates an outbound connection, the firewall maintains this address mapping. The address mapping structure is called an xlate, and remains in memory for a period of time. During this time, outside hosts can initiate connections to the inside host using the translated address from the pool, if allowed by the access list. Normally, outside-to-inside connections require a static translation so that the inside host always uses the same IP address.

- Service—Shows the service or protocol specified by the rule.
- Action—The action that applies to the rule, either Permit or Deny.
- Logging—If you enable logging for the access list, this column shows the logging level and the interval in seconds between log messages.
- Time—Displays the time range during which the rule is applied.
- Description—Shows the description you entered when you added the rule. An implicit rule includes the following description: “Implicit outbound rule.”
- Addresses—Tab that lets you add, edit, delete, or find IP names or network object groups. IP address objects are automatically created based on source and destination entries during rule creation so that they can easily be selected in the creation of subsequent rules. They cannot be added, edited, or deleted manually.
- Services—Tab that lets you add, edit, delete, or find services.
- Time Ranges—Tab that lets you add, edit, or delete time ranges.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Rule Queries

Configuration > Security Policy > Rule Queries

The Rule Queries dialog box lets you manage named rule queries that you can use in the Filter field when searching for Rules.

Fields

- Add—Adds a rule query.
- Edit—Edits a rule query.
- Delete—Deletes a rule query.
- Name—Lists the names of the rule queries.
- Description—Lists the descriptions of the rule queries.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

New/Edit Rule Query

Configuration > Security Policy > Rule Queries > New/Edit Rule Query

The New/Edit Rule Query dialog box lets you add or edit a named rule query that you can use in the Filter field when searching for Rules.

Fields

- Name—Enter a name for this rule query.
- Description—Enter a description for this rule query.
- Match Criteria—This area lists the criteria you want to filter on.
  - any of the following criteria—Sets the rule query to match any of the listed criteria.
  - all of the following criteria—Sets the rule query to match all of the listed criteria.
  - Field—Lists the type of criteria. For example, an interface or source.
  - Value—Lists the value of the criteria, for example, “inside.”
  - Remove—Removes the selected criteria.
- Define New Criteria—This area lets you define new criteria to add to the match criteria.
  - Field—Choose a type of criteria, including Interface, Source, Destination, Service, Action, or another Rule Query to be nested in this rule query.
  - Value—Enter a value to search on. For the Interface type, this field becomes a drop-down list so you can choose an interface name. For the Action type, the drop-down list includes Permit and Deny. For the Rule Query type, the drop-down list includes all defined rule queries. The Source and Destination types accept an IP address. You can type one manually, or browse for one by clicking the **...** button and launching the Browse Address dialog box. The Service type accepts a TCP, UDP, TCP-UDP, ICMP, or IP protocol type. You can type one manually, or browse for one by clicking the **...** button and launching the Browse Service Groups dialog box.
  - Add—Adds the criteria to the Match Criteria table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Access Rule

Configuration > Security Policy > Access Rules > Add/Edit Access Rule

The Add/Edit Rule dialog box lets you create a new rule, or modify an existing rule.

Fields

- Interface—Specifies the interface to which the rule applies.
- Action—Determines the action type of the new rule. Select either permit or deny.
  - Permit—Permits all matching traffic.
  - Deny—Denies all matching traffic.
- Direction—Determines which direction of traffic the rule is applied.
  - Incoming—Selects incoming traffic to the source interface.
  - Outgoing—Selects outgoing traffic from the destination interface.
- Source Type—Specifies the IP address, network object group, interface IP, or any, from which traffic is permitted or denied to the destination specified in the Destination Type field.
  - IP Address—Specifies the IP address from which traffic is permitted or denied to the destination specified in the Destination Type field.

    IP address—Specifies the IP address.

    ...—Lets you select, add, edit, delete, or find an existing IP address object, IP name, network object group, or all.

    Netmask—Specifies the netmask.
  - Network Object Group—Specifies the network object group from which traffic is permitted or denied to the destination specified in the Destination Type field.

    Group Name—Specifies the network object group name.

    ...—Lets you select, add, edit, delete, or find an existing IP address object, IP name, network object group, or all.
  - Interface IP—Specifies the interface IP, or any, from which traffic is permitted or denied to the destination specified in the Destination Type field.
  - Interface—Specifies the interface.
- Destination Type—Specifies the IP address, network object group, interface IP, or any, to which traffic is permitted or denied from the source specified in the Source Type field.
  - IP Address—Specifies the IP address to which traffic is permitted or denied from the destination specified in the Source Type field.

    IP address—Specifies the IP address.

    ...—Lets you select, add, edit, delete, or find an existing IP address object, IP name, network object group, or all.

    Netmask—Specifies the netmask.
  - Network Object Group—Specifies the network object group to which traffic is permitted or denied from the source specified in the Source Type field.

    Group Name—Specifies the network object group name.

    ...—Lets you select, add, edit, delete, or find an existing IP address object, IP name, network object group, or all.

  - Interface IP—Specifies the interface IP, or any, to which traffic is permitted or denied from the source specified in the Source Type field.
  - Interface—Specifies the interface.
- Protocol and Service: TCP and UDP—Selects the TCP/UDP protocol for the rule. The Source Port and Destination Port areas allow you to specify the ports that the access list uses to match packets.
  - Service—Choose this option to specify a port number, a range of ports, or a well-known service name from a list of services, such as HTTP or FTP. The operator drop-down list specifies how the access list matches the port. Choose one of the following operators:

    = —Equals the port number.

    not = —Does not equal the port number.

    > —Greater than the port number.

    < —Less than the port number.

    range—Equal to one of the port numbers in the range.
  - Group—Choose this option to specify a service group from the Service Group drop-down list,

    The browse button displays the Browse Source Port dialog box, which lets you select, add, edit, delete or find a source type from a preconfigured list.
- Protocol and Service: IP—Specifies the IP protocol for the rule in the IP protocol field.
- Protocol and Service: ICMP—Specifies the ICMP type for the rule in the ICMP type field or the ICMP group.
  - The browse button displays the Browse ICMP dialog box, which lets you select, add, edit, delete or find a source type from a preconfigured list.
- Rule Flow Diagram—Shows the networks, type of traffic, interface name, direction of flow, and action.
- Options—Enables logging for the access list and sets logging options. Logging options:
  - Use default logging behavior.
  - Enable logging for the rule. Sets the level and interval for permit and deny logging. See Log Options for more information.

    Syslog Level—Specifies emergencies, alerts, critical, errors, warnings, notifications, informational, or debugging.

    Log Interval—Specifies the interval for logging.
  - Disable logging for the rule.
  - Time Range—Select a time range defined for this rule from the drop-down list.

    The browse button displays the Browse Time Range dialog box, which lets you select, add, edit, or delete a time range from a preconfigured list.
  - Description—(Optional) Enter a description of the access rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Manage Service Groups

Configuration > Security Policy > Access Rules > Add/Edit Access Rule > Manage Service Groups

The Manage Service Groups dialog box lets you associate multiple TCP, UDP, or TCP-UDP services (ports) in a named group. You can then use the service group in an access or IPSec rule, a conduit, or other functions within ASDM and the CLI.

The term *service* refers to higher layer protocols associated with application level services having well known port numbers and "literal" names such as ftp, telnet, and smtp.

The security appliance permits the following TCP literal names:

bgp, chargen, cmd, daytime, discard, domain, echo, exec, finger, ftp, ftp-data, gopher, h323, hostname, http, ident, irc, klogin, kshell, lpd, nntp, pop2, pop3, pptp, smtp, sqlnet, sunrpc, tacacs, talk, telnet, time, uucp, whois, www.

The Name of a service group must be unique to all four types of object groups. For example, a service group and a network group may not share the same name.

Multiple service groups can be nested into a "group of groups" and used the same as a single group. When a service object group is deleted, it is removed from all service object groups where it is used.

If a service group is used in an access rule, do not remove it. A service group used in an access rule cannot be made empty.

Fields

- TCP—Select this option to add TCP services or port numbers to an object group.
- UDP—Select this option to add UDP services or port numbers to an object group.
- TCP-UDP—Select this option to add services or port numbers that are common to TCP and UDP to an object group.
- Service Group table—This table contains a descriptive name for each service object group. To modify or delete a group on this list, select the group and click Edit or Delete. To add a new group to this list, click Add.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Service Group

Configuration > Security Policy > Access Rules > Add/Edit Access Rule > Manage Service Groups > Add/Edit Service Group

The Add/Edit Service Group dialog box lets you manage a group of TCP/UDP services/ports.

Fields

- Service Group Name—Specifies the name of the service group. The name must be unique for all object groups. A service group name cannot share a name with a network group.
- Description—Specifies a description of the service group.
- Service—Lets you select services for the service group from a predefined drop-down list.
- Range/Port #—Lets you specify a range of ports for the service group.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Advanced Access Rule Configuration

Configuration > Security Policy > Access Rules > Advanced Access Rule Configuration

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

When you enable logging, if a packet matches the ACE, the security appliance creates a flow entry to track the number of packets received within a specific interval (see Log Options). The security appliance 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. At the end of each interval, the security appliance resets the hit count to 0. If no packets match the ACE during an interval, the security appliance 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 security appliance 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 security appliance does not create a new deny flow until the existing flows expire. If someone initiates a denial of service attack, the security appliance 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 control entry (also known as a rule) for the access list. See Log Options for more information.

Fields

- Maximum Deny-flows—The maximum number of deny flows permitted before the security appliance 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 list, the access list provided by a RADIUS server replaces the access list configured on that interface. If the per user override feature is disabled, the access list provided by the RADIUS server is combined with the access list configured on that interface. If the inbound access list is not configured for the interface, per user override cannot be configured.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Log Options

Configuration > Security Policy > Access Rules > Add/Edit Access Rule > Log Options (You can get to this dialog box through multiple paths.)

The Log Options dialog box lets you set logging options for each access control entry (also called a rule) for an access control list. Conduits and outbound lists do not support logging. See Advanced Access Rule Configuration to set global logging options.

This dialog box lets you use the older logging mechanism (only denied traffic is logged), to use the newer logging mechanism (permitted and denied traffic is logged, along with additional information such as how many packet hits), or to disable logging.

The Log option consumes a certain amount of memory when enabled. To help control the risk of a potential Denial of Service attack, you can configure the Maximum Deny-flow setting by choosing Advanced in the Access Rules window.

Fields

- Use default logging behavior—Uses the older access list logging mechanism: the security appliance logs system log message number 106023 when a packet is denied. Use this option to return to the default setting.
- Enable logging for the rule—Enables the newer access list logging mechanism: the security appliance logs system log message number 106100 when a packet matches the ACE (either permit or deny).

  If a packet matches the ACE, the security appliance creates a flow entry to track the number of packets received within a specific interval (see the Logging Interval field that follows). The security appliance 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. At the end of each interval, the security appliance resets the hit count to 0. If no packets match the ACE during an interval, the security appliance deletes the flow entry.

  – Logging Level—Selects the level of logging messages to be sent to the syslog server from this drop-down list. Levels are defined as follows:

Emergency (level 0)—The security appliance does not use this level.

Alert (level 1, immediate action needed)

Critical (level 2, critical condition)

Error (level 3, error condition)

Warning (level 4, warning condition)

Notification (level 5, normal but significant condition)

Informational (level 6, informational message only)

Debugging (level 7, appears during debugging only)

- Logging Interval—Sets the amount of time in seconds (1-600) the security appliance waits before sending the flow statistics to the syslog. This setting also serves as the timeout value for deleting a flow if no packets match the ACE. The default is 300 seconds.

- Disable logging for the rule—Disables all logging for the ACE.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

CHAPTER 18

Configuring EtherType Rules

Ethertype Rules (Transparent Mode Only)

Configuration > Security Policy > Ethertype Rules

The EtherType Rules window shows access rules based on packet EtherTypes. EtherType rules are used to configure non-IP related traffic policies through the security appliance 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.

Fields

- Add—Adds a new EtherType rule. Choose the type of rule you want to add from the drop-down list.
- Edit—Edits an EtherType rule.
- Delete—Deletes an EtherType rule.
- Move Up—Moves a rule up. Rules are assessed in the order they appear in this table, so the order can matter if you have overlapping rules.
- Move Down—Moves a rule down.
- Cut—Cuts a rule.
- Copy—Copies the parameters of a rule so you can start a new rule with the same parameters using the Paste button.
- Paste—Opens an Add/Edit Rule dialog box with the copied or cut parameters of the rule prefilled. You can then make any modifications and add it to the table. The Paste button adds the rule above the selected rule. The Paste After item, available from the Paste drop-down list, adds the rule after the selected rule.

The following description summarizes the columns in the EtherType Rules table. You can edit the contents of these columns by double-clicking on a table cell. Double-clicking on a column header sorts the table in ascending alphanumeric order, using the selected column as the sort key. If you right-click a rule, you see all of the options represented by the buttons above, as well as Insert and Insert After items. These items either insert a new rule before the selected rule (Insert) or after the selected rule (Insert After.)

- No—Indicates the order of evaluation for the rule.
- Action—Permit or deny action for this rule.
- Ethervalue—EtherType value: IPX, BPDU, MPLS-Unicast, MPLS-Multicast, or a 16-bit hexadecimal value between 0x600 (1536) and 0xffff by which an EtherType can be identified.
- Interface—Interface to which the rule is applied.

- Direction Applied—Direction for this rule: incoming traffic or outgoing traffic.
- Description—Optional text description of the rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
—	•	•	•	—

Add/Edit EtherType Rule

Configuration > Security Policy > Ethertype Rules > Add/Edit Ethertype Rules

The Add/Edit EtherType Rules dialog box lets you add or edit an EtherType rule.

Fields

- Action—Permit or deny action for this rule.
- Interface—Interface name for this rule.
- Apply rule to—Direction for this rule: incoming traffic or outgoing traffic.
- Ethervalue—EtherType value: BPDU, IPX, MPLS-Unicast, MPLS-Multicast, any (any value between 0x600 and 0xffff), or a 16-bit hexadecimal value between 0x600 (1536) and 0xffff by which an EtherType can be identified.
- Description—Optional text description of the rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
—	•	•	•	—

Configuring AAA Rules

This chapter describes how to enable AAA (pronounced “triple A”) for network access.

For information about AAA for management access, see the “AAA Access” section on page 11-1

This chapter includes the following sections:

AAA Performance

The security appliance uses “cut-through proxy” to significantly improve performance compared to a traditional proxy server. The performance of a traditional proxy server suffers because it analyzes every packet at the application layer of the OSI model. The security appliance cut-through proxy challenges a user initially at the application layer and then authenticates against standard AAA servers or the local database. After the security appliance authenticates the user, it shifts the session flow, and all traffic flows directly and quickly between the source and destination while maintaining session state information.

Configuring AAA Rules

This section describes how to configure AAA rules, and includes the following topics:

AAA Rules

Configuration > Security Policy > AAA Rules

The Security Policy pane shows your network security policy expressed in rules. This window includes tabs for AAA Rules, as well as for other rules. This topic describes AAA Rules. For an overview of AAA services, see Chapter 10, "Configuring AAA Servers."

When you choose the **AAA Rules** tab, you can define authentication, authorization, or accounting (AAA) rules, as well as MAC exempt rules. AAA tells the security appliance who the user is, what the user can do, and what the user did. You can use authentication alone, or with authorization. Authorization always requires authentication. For example, if you authenticate outside users who access any server on the inside network, then authentication alone is adequate. However, if you want to limit the inside servers that a particular user accesses, you can configure an authorization server to specify which servers and services that user is allowed to access.

AAA provides a greater level of protection and control for user access than using access lists alone. For example, you can create an access list allowing all outside users to access a server on the DMZ network. But if you want only registered users to Telnet to the server, you can configure AAA to allow only authenticated and/or authorized users to make it past the security appliance. If the server also has its own authentication and authorization, the user enters a second set of user name and password (in the case of FTP, the user must enter both usernames and passwords separated by an at sign (@)).

Each AAA rule identifies the following characteristics for matching traffic:

- The source and destination network
- The action (authentication, authorization, or accounting; a rule can also exempt a MAC address from AAA)
- The AAA server group
- The service type (for example, Telnet or FTP)

Restrictions

ASDM does not support a mixed configuration of AAA rules that:

- Specify the source and destination addresses
- Match access lists for the source and destination addresses

  If your configuration already contains AAA rules, then you can add only AAA rules of the same kind. If you have not configured any AAA rules, then ASDM allows you to add only rules that match access lists. To convert your rules to match access lists, you must delete all of your AAA rules in ASDM, then re-add them (with no rules configured, ASDM defaults to access list mode). In ASDM, the configuration of AAA rules is the same in both modes.
- For FTP authentication, the user must enter the name and password in the following format:

  security appliance_*name@ftp_name*
  security appliance_*password@ftp_password*
- The security appliance forwards the FTP name and password to the FTP server after successful authentication on the security appliance. Other services such as Telnet and HTTP (if configured for authentication) require you to enter a second name and password at the destination server prompt.
- Some services are not reliably authenticated, such as mail or SMTP. If you specify that *all* services need to be authenticated, then the user must first authenticate with Telnet, FTP, HTTP, or HTTPS (or another service that reliably provides an authentication prompt), and then use the other services.
- AAA authorization rules support TACACS+ servers, but not other servers. However, you can use the local database to authorize users for security appliance commands.
- AAA accounting rules are not supported using the local database as the AAA Server Group.

Prerequisites

1. Define each host or server in the Configuration > Features > Properties > AAA Setup > AAA Server Groups pane.
2. Add users to the local database. (See Configuration > Features > Properties > Administration > User Accounts.)
3. Be sure that users can access the specified network (by an Access Rules if required).
4. Set up the AAA server correctly.

Fields

- Add—Adds a new AAA rule. Choose the type of rule you want to add from the drop-down list.
- Edit—Edits an AAA rule.
- Delete—Deletes a AAA rule.
- Move Up—Moves a rule up. Rules are assesed in the order they appear in this table, so the order can matter if you have overlapping rules.
- Move Down—Moves a rule down.
- Cut—Cuts a rule.
- Copy—Copies a rule's parameters so you can start a new rule with the same parameters using the Paste button.
- Paste—Opens an Add/Edit Rule dialog box with the copied or cut rule's paramaters prefilled. You can then make any modifications and add it to the table. The Paste button adds the rule above the selected rule. The Paste After item, available from the Paste drop-down list, adds the rule after the selected rule.
- Find—Filters the display to show only matching rules. Clicking **Find** opens the Filter field. Click **Find** again to hide the Filter field.
  - Filter drop-down list—Choose the criteria to filter on, either Interface, Source, Destination, Service, Action, or Rule Query. A rule query is a collection of multiple criteria that you can save and use repeatedly.
  - Filter field—For the Interface type, this field becomes a drop-down list so you can choose an interface name, or **All Interfaces**. For the Action type, the drop-down list includes Permit and Deny. For the Rule Query type, the drop-down list includes all defined rule queries. The Source and Destination types accept an IP address. You can type one manually, or browse for one by clicking the **...** button and launching the Browse Address dialog box. The Service type accepts a TCP, UDP, TCP-UDP, ICMP, or IP protocol type. You can type one manually, or browse for one by clicking the **...** button and launching the Browse Service Groups dialog box.
  - Filter—Runs the filter.
  - Clear—Clears the Filter field.
  - Rule Query—Opens the Rule Queries dialog box so you can manage named rule queries.
- Show Rule Flow Diagram—Shows the Rule Flow Diagram area under the rule table. This diagram shows the networks, type of traffic, interface name, direction of flow, and action (for example, Authenticate or Do Not Authenticate).
- Packet Trace—Opens the Packet Tracer tool with the paramters pre-filled with the characteristics of the selected rule.

The following description summarizes the columns in the AAA Rules table. You can edit the contents of these columns by double-clicking on a table cell. Double-clicking on a column header sorts the table in ascending alphanumeric order, using the selected column as the sort key. If you right-click a rule, you see all of the options represented by the buttons above, as well as Insert and Insert After items. These items either insert a new rule before the selected rule (Insert) or after the selected rule (Insert After.)

- No—Indicates the order of evaluation for the rule.
- Enabled—Indicates whether the rule is enabled or disabled.
- Action—Specifies the type of AAA rule.
- Source—Lists the IP addresses that are subject to AAA when traffic is sent to the IP addresses listed in the Destination column.
- Destination—Lists the IP addresses that are subject to AAA when traffic is sent from the IP addresses listed in the Source column.
- Service—Shows the service or protocol specified by the rule.
- Action—Shows the action specified by the rule, including Authenticate, Do Not Authenticate, Authorize, Do Not Authorize, and so on.
- Server Group—Specifies the AAA Server Group tag. Configure AAA server groups in Properties > AAA Setup > AAA Server Groups. To create new AAA rules, a server group must exist and have one or more servers in it.
- Time—Specifies the name of the time range in effect for this rule.
- Description—The description you entered when you added the rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Authentication Rule

The security appliance lets you configure network access authentication using AAA servers or the local database.

A user at a given IP address only needs to authenticate one time for all rules and types, until the authentication session expires. (See Timeouts for timeout values.) For example, if you configure the security appliance to authenticate Telnet and FTP, and a user first successfully authenticates for Telnet, then as long as the authentication session exists, the user does not also have to authenticate for FTP.

Although you can configure the security appliance to require authentication for network access to any protocol or service, users can authenticate directly with HTTP(S), Telnet, or FTP only. A user must first authenticate with one of these services before the security appliance allows other traffic requiring authentication.

If you do not want to allow HTTP(S), Telnet, or FTP through the security appliance but want to authenticate other types of traffic, you can configure virtual Telnet. With virtual Telnet, the user Telnets to a given IP address configured on the security appliance and the security appliance provides a Telnet prompt.

For Telnet, HTTP(S), and FTP, the security appliance generates an authentication prompt. If the destination server also has its own authentication, the user enters another username and password.

For HTTP authentication, the security appliance checks local ports when static NAT is configured. If it detects traffic destined for local port 80, regardless of the global port, the security appliance intercepts the HTTP connection and enforces authentication.

For example, assume that outside TCP port 889 is translated to port 80 (www) and that any relevant access lists permit the traffic.

Then when users try to access 10.48.66.155 on port 889, the security appliance intercepts the traffic and enforces HTTP authentication. Users see the HTTP authentication page in their web browsers before the security appliance allows HTTP connection to complete.

If the local port is different than port 80, then users do not see the authentication page. Instead, the security appliance sends to the web browser an error message indicating that the user must be authenticated prior using the requested service.

Note If you use HTTP authentication without using secure HTTP client authentication (see Configuring Advanced AAA Features), the username and password are sent in clear text to the destination web server, and not just to the AAA server. For example, if you authenticate inside users when they access outside web servers, anyone on the outside can learn valid usernames and passwords. We recommend that you use secure HTTP client authentication whenever you enable HTTP authentication.

For FTP, a user has the option of entering the security appliance username followed by an at sign (@) and then the FTP username (name1@name2). For the password, the user enters the security appliance password followed by an at sign (@) and then the FTP password (password1@password2). For example, enter the following text.

```
name> jamiec@jchrichton
password> letmein@he110
```

This feature is useful when you have cascaded firewalls that require multiple logins. You can separate several names and passwords by multiple at signs (@).

Fields

Interface and Action—Choose the interface, action, and AAA server group.

- Interface—Choose the interface on which to apply this rule.
- Action—Choose **Authenticate** or **Do not Authenticate**.
- AAA Server Group—Choose a AAA server group or the local database. You must add the server group in Properties > AAA Setup > AAA Server Groups.
- Add Server/User—Click this button to add a server to the selected AAA server group, or a user to the local database.

Source—Specify the source address for traffic you want to authenticate.

- Type—Choose the type of address you want to use, including any, IP address, Network Object Group, or Interface IP.

  If you choose **IP address**, the you see the following fields:

- IP Address—Enter it manually or click the **...** button to choose from the Browse Address dialog box.
- Netmask—Choose a subnet mask from the drop-down list.

If you choose **Network Object Group**, you see the following field:

- Group Name—Choose a group name from the drop-down list, or click the **...** button to bring up the Browse Address dialog box. From the Browse Address dialog box, you can add a network object group.

If you choose **Interface IP**, you see the following field:

- Interface—Choose an interface from the drop-down list.

Destination—Specify the destination address for traffic you want to authenticate.

- Type—Choose the type of address you want to use, including any, IP address, Network Object Group, or Interface IP.

  If you choose **IP address**, the you see the following fields:

  - IP Address—Enter it manually or click the **...** button to choose from the Browse Address dialog box.
  - Netmask—Choose a subnet mask from the drop-down list.

  If you choose **Network Object Group**, you see the following field:

  - Group Name—Choose a group name from the drop-down list, or click the **...** button to bring up the Browse Address dialog box. From the Browse Address dialog box, you can add a network object group.

  If you choose **Interface IP**, you see the following field:

  - Interface—Choose an interface from the drop-down list.

Protocol and Service—Specify the port or protocol for traffic you want to authenticate.

- Protocol—Choose the protocol for the traffic, either tcp, udp, ip, icmp, or other.

  If you choose **tcp** or **udp**, then you see the following fields:

  - Source Port—Set the source port for the traffic you want to authenticate.

    Service—Click this radio button to enter a port or range of ports. Choose an operator from the drop-down list, including = (equal), != (not equal), > (greater than), < (less than), and range. Either type a number, or choose a well-known port name from the drop-down list. For a range, you must specify numbers.

    Group—Click this radio button to specify a service group, created on Service Groups. Choose a group name from the drop-down list, or click the **...** button to bring up the Browse Service Groups dialog box. From the Browse Service Groups dialog box, you can add a service group.

  - Destination Port—Set the destination port for the traffic you want to authenticate.

    Service—Click this radio button to enter a port or range of ports. Choose an operator from the drop-down list, including = (equal), != (not equal), > (greater than), < (less than), and range. Either type a number, or choose a well-known port name from the drop-down list. For a range, you must specify numbers.

    Group—Click this radio button to specify a service group, created on Service Groups. Choose a group name from the drop-down list, or click the **...** button to bring up the Browse Service Groups dialog box. From the Browse Service Groups dialog box, you can add a service group.

  If you choose **icmp**, then you see the following fields:

- ICMP Type—Click this radio button to enter an ICMP type. Either type a number, or choose a well-known type from the drop-down list.
- ICMP Group—Click this radio button to specify a service group, created on Service Groups. Choose a group name from the drop-down list, or click the **...** button to bring up the Browse Service Groups dialog box. From the Browse Service Groups dialog box, you can add a service group.

If you choose **other**, then you see the following fields:

- Protocol—Click this radio button to enter an IP protocol type. Either type a number, or choose a well-known type from the drop-down list.
- Protocol Group—Click this radio button to specify a service group, created on Service Groups. Choose a group name from the drop-down list, or click the **...** button to bring up the Browse Service Groups dialog box. From the Browse Service Groups dialog box, you can add a service group.

Rule Flow Diagram—Shows the Rule Flow Diagram for this rule. This diagram shows the networks, type of traffic, interface name, direction of flow, and action (for example, Authenticate or Do Not Authenticate).

Options—Set options for this rule.

- Time Range—Choose the name of an existing time range from the drop-down list. A time range enables a rule only during the specified times. Create a time range on Configuring Time Ranges.
- Description—Enter a description for this rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Authorization Rule

You can configure the security appliance to perform network access authorization with TACACS+.

Note When you configure the security appliance to authenticate users for network access using RADIUS, you are also implicitly enabling RADIUS authorizations. RADIUS authorization does not require a separate authorization rule, like TACACS+. See the "Configuring a RADIUS Server for Authorization" section on page 19-15 for more information about using RADIUS for authorization.

Authentication and authorization rules are independent; however, any unauthenticated traffic matched by an authorization rule will be denied. For authorization to succeed, a user must first authenticate with the security appliance. Because a user at a given IP address only needs to authenticate one time for all rules and types, if the authentication session has not expired, authorization can occur even if the traffic is matched by an authentication statement.

After a user authenticates, the security appliance checks the authorization rules for matching traffic. If the traffic matches the authorization rule, the security appliance sends the username to the TACACS+ server. The TACACS+ server responds to the security appliance with a permit or a deny for that traffic, based on the user profile. The security appliance enforces the authorization rule in the response.

See the documentation for your TACACS+ server for information about configuring network access authorizations for a user.

Fields

Interface and Action—Choose the interface, action, and AAA server group.

- Interface—Choose the interface on which to apply this rule.
- Action—Choose **Authorize** or **Do not Authorize**.
- AAA Server Group—Choose a AAA server group or the local database. You must add the server group in Properties > AAA Setup > AAA Server Groups.
- Add Server/User—Click this button to add a server to the selected AAA server group, or a user to the local database.

Source—Specify the source address for traffic you want to authorize.

- Type—Choose the type of address you want to use, including any, IP address, Network Object Group, or Interface IP.

  If you choose **IP address**, the you see the following fields:

  – IP Address—Enter it manually or click the **...** button to choose from the Browse Address dialog box.

  – Netmask—Choose a subnet mask from the drop-down list.

  If you choose **Network Object Group**, you see the following field:

  – Group Name—Choose a group name from the drop-down list, or click the **...** button to bring up the Browse Address dialog box. From the Browse Address dialog box, you can add a network object group.

  If you choose **Interface IP**, you see the following field:

  – Interface—Choose an interface from the drop-down list.

Destination—Specify the destination address for traffic you want to authorize.

- Type—Choose the type of address you want to use, including any, IP address, Network Object Group, or Interface IP.

  If you choose **IP address**, the you see the following fields:

  – IP Address—Enter it manually or click the **...** button to choose from the Browse Address dialog box.

  – Netmask—Choose a subnet mask from the drop-down list.

  If you choose **Network Object Group**, you see the following field:

  – Group Name—Choose a group name from the drop-down list, or click the **...** button to bring up the Browse Address dialog box. From the Browse Address dialog box, you can add a network object group.

  If you choose **Interface IP**, you see the following field:

  – Interface—Choose an interface from the drop-down list.

Protocol and Service—Specify the port or protocol for traffic you want to authorize.

- Protocol—Choose the protocol for the traffic, either tcp, udp, ip, icmp, or other.

  If you choose **tcp** or **udp**, then you see the following fields:

  - Source Port—Set the source port for the traffic you want to authorize.

    Service—Click this radio button to enter a port or range of ports. Choose an operator from the drop-down list, including = (equal), != (not equal), > (greater than), < (less than), and range. Either type a number, or choose a well-known port name from the drop-down list. For a range, you must specify numbers.

    Group—Click this radio button to specify a service group, created on Service Groups. Choose a group name from the drop-down list, or click the **...** button to bring up the Browse Service Groups dialog box. From the Browse Service Groups dialog box, you can add a service group.

  - Destination Port—Set the destination port for the traffic you want to authorize.

    Service—Click this radio button to enter a port or range of ports. Choose an operator from the drop-down list, including = (equal), != (not equal), > (greater than), < (less than), and range. Either type a number, or choose a well-known port name from the drop-down list. For a range, you must specify numbers.

    Group—Click this radio button to specify a service group, created on Service Groups. Choose a group name from the drop-down list, or click the **...** button to bring up the Browse Service Groups dialog box. From the Browse Service Groups dialog box, you can add a service group.

  If you choose **icmp**, then you see the following fields:

  - ICMP Type—Click this radio button to enter an ICMP type. Either type a number, or choose a well-known type from the drop-down list.
  - ICMP Group—Click this radio button to specify a service group, created on Service Groups. Choose a group name from the drop-down list, or click the **...** button to bring up the Browse Service Groups dialog box. From the Browse Service Groups dialog box, you can add a service group.

  If you choose **other**, then you see the following fields:

  - Protocol—Click this radio button to enter an IP protocol type. Either type a number, or choose a well-known type from the drop-down list.
  - Protocol Group—Click this radio button to specify a service group, created on Service Groups. Choose a group name from the drop-down list, or click the **...** button to bring up the Browse Service Groups dialog box. From the Browse Service Groups dialog box, you can add a service group.

Rule Flow Diagram—Shows the Rule Flow Diagram for this rule. This diagram shows the networks, type of traffic, interface name, direction of flow, and action (for example, authorize or Do Not Authorize).

Options—Set options for this rule.

- Time Range—Choose the name of an existing time range from the drop-down list. A time range enables a rule only during the specified times. Create a time range on Configuring Time Ranges.
- Description—Enter a description for this rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Accounting Rule

The security appliance can send accounting information to a RADIUS or TACACS+ server about any TCP or UDP traffic that passes through the security appliance. If that traffic is also authenticated, then the AAA server can maintain accounting information by username. If the traffic is not authenticated, the AAA server can maintain accounting information by IP address. Accounting information includes when sessions start and stop, username, the number of bytes that pass through the security appliance for the session, the service used, and the duration of each session.

Fields

Interface and Action—Choose the interface, action, and AAA server group.

- Interface—Choose the interface on which to apply this rule.
- Action—Choose **Account** or **Do not Account**.
- AAA Server Group—Choose a AAA server group or the local database. You must add the server group in Properties > AAA Setup > AAA Server Groups.
- Add Server/User—Click this button to add a server to the selected AAA server group, or a user to the local database.

Source—Specify the source address for traffic you want to authenticate.

- Type—Choose the type of address you want to use, including any, IP address, Network Object Group, or Interface IP.

  If you choose **IP address**, the you see the following fields:

  – IP Address—Enter it manually or click the **...** button to choose from the Browse Address dialog box.

  – Netmask—Choose a subnet mask from the drop-down list.

  If you choose **Network Object Group**, you see the following field:

  – Group Name—Choose a group name from the drop-down list, or click the **...** button to bring up the Browse Address dialog box. From the Browse Address dialog box, you can add a network object group.

  If you choose **Interface IP**, you see the following field:

  – Interface—Choose an interface from the drop-down list.

Destination—Specify the destination address for traffic you want to account.

- Type—Choose the type of address you want to use, including any, IP address, Network Object Group, or Interface IP.

  If you choose **IP address**, the you see the following fields:

  – IP Address—Enter it manually or click the **...** button to choose from the Browse Address dialog box.

  - Netmask—Choose a subnet mask from the drop-down list.

If you choose **Network Object Group**, you see the following field:

  - Group Name—Choose a group name from the drop-down list, or click the **...** button to bring up the Browse Address dialog box. From the Browse Address dialog box, you can add a network object group.

If you choose **Interface IP**, you see the following field:

  - Interface—Choose an interface from the drop-down list.

Protocol and Service—Specify the port or protocol for traffic you want to account.

- Protocol—Choose the protocol for the traffic, either tcp or udp.
  - Source Port—Set the source port for the traffic you want to account.

    Service—Click this radio button to enter a port or range of ports. Choose an operator from the drop-down list, including = (equal), != (not equal), > (greater than), < (less than), and range. Either type a number, or choose a well-known port name from the drop-down list. For a range, you must specify numbers.

    Group—Click this radio button to specify a service group, created on Service Groups. Choose a group name from the drop-down list, or click the **...** button to bring up the Browse Service Groups dialog box. From the Browse Service Groups dialog box, you can add a service group.
  - Destination Port—Set the destination port for the traffic you want to account.

    Service—Click this radio button to enter a port or range of ports. Choose an operator from the drop-down list, including = (equal), != (not equal), > (greater than), < (less than), and range. Either type a number, or choose a well-known port name from the drop-down list. For a range, you must specify numbers.

    Group—Click this radio button to specify a service group, created on Service Groups. Choose a group name from the drop-down list, or click the **...** button to bring up the Browse Service Groups dialog box. From the Browse Service Groups dialog box, you can add a service group.

Rule Flow Diagram—Shows the Rule Flow Diagram for this rule. This diagram shows the networks, type of traffic, interface name, direction of flow, and action (for example, Account or Do Not Account).

Options—Set options for this rule.

- Time Range—Choose the name of an existing time range from the drop-down list. A time range enables a rule only during the specified times. Create a time range on Configuring Time Ranges.
- Description—Enter a description for this rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit MAC Exempt Rule

The security appliance can exempt from authentication and authorization any traffic from specific MAC addresses.

For example, if the security appliance authenticates TCP traffic originating on a particular network but you want to allow unauthenticated TCP connections from a specific server, you would use a MAC exempt rule to exempt from authentication and authorization any traffic from the server specified by the rule.

The order of entries matters, because the packet uses the first entry it matches, as opposed to a best match scenario. If you have a permit entry, and you want to deny an address that is allowed by the permit entry, be sure to enter the deny entry before the permit entry.

Fields

- Action—Choose **MAC Exempt** or **No MAC Exempt**. The MAC Exempt option allows traffic from the MAC address without having to authenticate or authorize. The No MAC Exempt option specifies a MAC address that is not exempt from authentication or authorization. You might need to add a deny entry if you permit a range of MAC addresses using a MAC address mask such as ffff.ffff.0000, and you want to force a MAC address in that range to be authenticated and authorized.
- MAC Address—Specifies the source MAC address in 12-digit hexadecimal form; that is, nnnn.nnnn.nnnn.
- MAC Mask—Specifies the portion of the MAC address that should be used for matching. For example, ffff.ffff.ffff matches the MAC address exactly. ffff.ffff.0000 matches only the first 8 digits.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Configuring Advanced AAA Features

Configuration > Security Policy > AAA Rules > Advanced AAA Configuration

The Advanced AAA Configuration dialog box enables secure HTTP, sets the Proxy Limit, and enables interactive authentication.

Fields

- Secure HTTP—Specifies whether to enable or disable Secure HTTP (HTTPS).
- Enable Secure HTTP—Enables secure HTTP authentication. Without securing HTTP authentication, usernames and passwords from the client to the security appliance are passed as clear text. By enabling this option, you enable the exchange of usernames and passwords between a web client and the security appliance with HTTPS. After enabling this feature, when a user requires authentication when using HTTP, the security appliance redirects the HTTP user to an HTTPS prompt. After you authenticate correctly, the security appliance redirects you to the original HTTP URL.

- Proxy Limit—Specifies Proxy Limit parameters.
  - Enable Proxy Limit—Limits the number of concurrent proxy connections allowed per user. The maximum is 128. If you do not enable this feature, no limit is imposed.
  - Proxy Limit— Specifies the number of concurrent proxy connections allowed. The range is 1 through 128. The default is 16.
- Interactive Authentication—Configures interactive authentication for HTTP and HTTPS traffic. The default is to use inline basic authentication. This area also configures direct authentication. See the "Adding an Interactive Authentication Rule" section on page 19-13 for detailed information about interactive authentication.
  - Interface—Shows the interface on which you enabled interactive authentication.
  - Protocol—Shows the protocol, HTTP or HTTPS.
  - Port—Shows the listening port.
  - Redirect—Shows whether you enabled redirection for through traffic. Without redirection, this rule only enables direct authentication.
  - Add—Adds an interactive authentication rule.
  - Edit—Edits an interactive authentication rule.
  - Delete—Deletes an interactive authentication rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Adding an Interactive Authentication Rule

By default for HTTP, the security appliance uses basic HTTP authentication. For HTTPS, the security appliance generates a similar custom login screen. Using the Configuration > Security Policy > AAA Rules > Advanced AAA Configuration > Add Interactive Authentication dialog box, you can configure the security appliance to redirect users to an internal web page where they can enter their username and password.

If you enable the redirect method of HTTP and HTTPS authentication, then you also automatically enable direct authentication with the security appliance. Direct authentication is useful if you do not want to allow HTTP, HTTPS, Telnet, or FTP through the security appliance but want to authenticate other types of traffic; a user can authenticate directly with the security appliance using HTTP or HTTPS before other traffic is allowed. You can configure direct authentication independently if you want to continue to use basic HTTP authentication for through traffic. To access the login page for direct authentication, enter one of the following URLs:

http://*interface_ip*[**:***port*]**/netaccess/connstatus.html**
https://*interface_ip*[**:***port*]**/netaccess/connstatus.html**

Redirection is an improvement over the basic method because it provides an improved user experience when authenticating, and an identical user experience for HTTP and HTTPS in both Easy VPN and firewall modes. It also supports authenticating directly with the security appliance.

You might want to continue to use basic HTTP authentication if: you do not want the security appliance to open listening ports; you use NAT on a router and you do not want to create a translation rule for the web page served by the security appliance; basic HTTP authentication might work better with your network. For example non-browser applications, like when a URL is embedded in email, might be more compatible with basic authentication.

To configure an interactive authentication rule, perform the following steps:

Step 1 From the Configuration > Security Policy > AAA Rules > Advanced AAA Configuration dialog box, click **Add**.

Step 2 From the Protocol menu, choose **HTTP** or **HTTPS**.

To enable listeners for both HTTP and HTTPS, you need to create two separate rules.

Step 3 From the Interface menu, choose the interface name on which you want to enable the listener.

Step 4 From the Port menu, either choose a common port or type the port number on which you want to listen.

The default is 80 for HTTP and 443 for HTTPS.

Step 5 To redirect through traffic to the listening port for authentication, check the **Redirect network users for authentication requests** check box.

If you do not check this check box, then only direct authentication is enabled.

Step 6 Click **OK**.

Fields

- Protocol—Sets the protocol for the interactive authentication rule, either HTTP or HTTPS.
- Interface—Sets the interface name on which you want to enable the listening port.
- Port—Sets the port number on which you want to listen. Choose a common port or type the port number. The default is 80 for HTTP and 443 for HTTPS.
- Redirect network users for authentication requests—Redirects through traffic to the listening port for authentication. If you do not check this check box, then only direct authentication is enabled.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Configuring a RADIUS Server for Authorization

When authentication succeeds, the RADIUS protocol returns user authorizations in the access-accept message sent by a RADIUS server.

When you configure the security appliance to authenticate users for network access, you are also implicitly enabling RADIUS authorizations; therefore, this section contains no information about configuring RADIUS authorization on the security appliance. It does provide information about how the security appliance handles access list information received from RADIUS servers.

You can configure a RADIUS server to download an access list to the security appliance or an access list name at the time of authentication. The user is authorized to do only what is permitted in the user-specific access list.

Note If you have used the **access-group** command to apply access lists to interfaces, be aware of the following effects of the **per-user-override** keyword on authorization by user-specific access lists:

- Without the **per-user-override** keyword, traffic for a user session must be permitted by both the interface access list and the user-specific access list.
- With the **per-user-override** keyword, the user-specific access list determines what is permitted.

For more information, see the **access-group** command entry in the *Cisco Security Appliance Command Reference*.

This section includes the following topics:

Configuring a RADIUS Server to Send Downloadable Access Control Lists

This section describes how to configure Cisco Secure ACS or a third-party RADIUS server, and includes the following topics:

About the Downloadable Access List Feature and Cisco Secure ACS

Downloadable access lists is the most scalable means of using Cisco Secure ACS to provide the appropriate access lists for each user. It provides the following capabilities:

- Unlimited access list size—Downloadable access lists are sent using as many RADIUS packets as required to transport the full access list from Cisco Secure ACS to the security appliance.
- Simplified and centralized management of access lists—Downloadable access lists enable you to write a set of access lists once and apply it to many user or group profiles and distribute it to many security appliances.

This approach is most useful when you have very large access list sets that you want to apply to more than one Cisco Secure ACS user or group; however, its ability to simplify Cisco Secure ACS user and group management makes it useful for access lists of any size.

The security appliance receives downloadable access lists from Cisco Secure ACS using the following process:

1. The security appliance sends a RADIUS authentication request packet for the user session.
2. If Cisco Secure ACS successfully authenticates the user, Cisco Secure ACS returns a RADIUS access-accept message that contains the internal name of the applicable downloadable access list. The Cisco IOS cisco-av-pair RADIUS VSA (vendor 9, attribute 1) contains the following attribute-value pair to identify the downloadable access list set:

   `ACS:CiscoSecure-Defined-ACL=`*`acl-set-name`*

   where *acl-set-name* is the internal name of the downloadable access list, which is a combination of the name assigned to the access list by the Cisco Secure ACS administrator and the date and time that the access list was last modified.
3. The security appliance examines the name of the downloadable access list and determines if it has previously received the named downloadable access list.
   - If the security appliance has previously received the named downloadable access list, communication with Cisco Secure ACS is complete and the security appliance applies the access list to the user session. Because the name of the downloadable access list includes the date and time it was last modified, matching the name sent by Cisco Secure ACS to the name of an access list previous downloaded means that the security appliance has the most recent version of the downloadable access list.
   - If the security appliance has not previously received the named downloadable access list, it may have an out-of-date version of the access list or it may not have downloaded any version of the access list. In either case, the security appliance issues a RADIUS authentication request using the downloadable access list name as the username in the RADIUS request and a null password attribute. In a cisco-av-pair RADIUS VSA, the request also includes the following attribute-value pairs:

     ```
     AAA:service=ip-admission
     AAA:event=acl-download
     ```

     In addition, the security appliance signs the request with the Message-Authenticator attribute (IETF RADIUS attribute 80).
4. Upon receipt of a RADIUS authentication request that has a username attribute containing the name of a downloadable access list, Cisco Secure ACS authenticates the request by checking the Message-Authenticator attribute. If the Message-Authenticator attribute is missing or incorrect, Cisco Secure ACS ignores the request. The presence of the Message-Authenticator attribute prevents malicious use of a downloadable access list name to gain unauthorized network access. The Message-Authenticator attribute and its use are defined in RFC 2869, RADIUS Extensions, available at http://www.ietf.org.
5. If the access list required is less than approximately 4 KB in length, Cisco Secure ACS responds with an access-accept message containing the access list. The largest access list that can fit in a single access-accept message is slightly less than 4 KB because some of the message must be other required attributes.

   Cisco Secure ACS sends the downloadable access list in a cisco-av-pair RADIUS VSA. The access list is formatted as a series of attribute-value pairs that each contain an ACE and are numbered serially:

   `ip:inacl#1=`*`ACE-1`*

```
ip:inacl#2=ACE-2
.
.
.
ip:inacl#n=ACE-n
```

An example of an attribute-value pair follows:

```
ip:inacl#1=permit tcp 10.1.0.0 255.0.0.0 10.0.0.0 255.0.0.0
```

6. If the access list required is more than approximately 4 KB in length, Cisco Secure ACS responds with an access-challenge message that contains a portion of the access list, formatted as described above, and an State attribute (IETF RADIUS attribute 24), which contains control data used by Cisco Secure ACS to track the progress of the download. Cisco Secure ACS fits as many complete attribute-value pairs into the cisco-av-pair RADIUS VSA as it can without exceeding the maximum RADIUS message size.

   The security appliance stores the portion of the access list received and responds with another access-request message containing the same attributes as the first request for the downloadable access list plus a copy of the State attribute received in the access-challenge message.

   This repeats until Cisco Secure ACS sends the last of the access list in an access-accept message.

Configuring Cisco Secure ACS for Downloadable Access Lists

You can configure downloadable access lists on Cisco Secure ACS as a shared profile component and then assign the access list to a group or to an individual user.

The access list definition consists of one or more security appliance commands that are similar to the extended **access-list** command, except without the following prefix:

access-list *acl_name* **extended**

The following example is a downloadable access list definition on Cisco Secure ACS version 3.3:

```
+-------------------------------------------+
| Shared profile Components                 |
|                                           |
|      Downloadable IP ACLs Content         |
|                                           |
| Name:    acs_ten_acl                      |
|                                           |
|      ACL Definitions                      |
|                                           |
| permit tcp any host 10.0.0.254            |
| permit udp any host 10.0.0.254            |
| permit icmp any host 10.0.0.254           |
| permit tcp any host 10.0.0.253            |
| permit udp any host 10.0.0.253            |
| permit icmp any host 10.0.0.253           |
| permit tcp any host 10.0.0.252            |
| permit udp any host 10.0.0.252            |
| permit icmp any host 10.0.0.252           |
| permit ip any any                         |
+-------------------------------------------+
```

For more information about creating downloadable access lists and associating them with users, see the user guide for your version of Cisco Secure ACS.

On the security appliance, the downloaded access list has the following name:

```
#ACSACL#-ip-acl_name-number
```

The *acl_name* argument is the name that is defined on Cisco Secure ACS (acs_ten_acl in the preceding example), and *number* is a unique version ID generated by Cisco Secure ACS.

The downloaded access list on the security appliance consists of the following lines:

```
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit tcp any host 10.0.0.254
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit udp any host 10.0.0.254
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit icmp any host 10.0.0.254
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit tcp any host 10.0.0.253
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit udp any host 10.0.0.253
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit icmp any host 10.0.0.253
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit tcp any host 10.0.0.252
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit udp any host 10.0.0.252
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit icmp any host 10.0.0.252
access-list #ACSACL#-ip-asa-acs_ten_acl-3b5385f7 permit ip any any
```

Configuring Any RADIUS Server for Downloadable Access Lists

You can configure any RADIUS server that supports Cisco IOS RADIUS VSAs to send user-specific access lists to the security appliance in a Cisco IOS RADIUS cisco-av-pair VSA (vendor 9, attribute 1).

In the cisco-av-pair VSA, configure one or more ACEs that are similar to the **access-list extended** command, except that you replace the following command prefix:

access-list *acl_name* **extended**

with the following text:

ip:inacl#*nnn***=**

The *nnn* argument is a number in the range from 0 to 999999999 that identifies the order of the command statement to be configured on the security appliance. If this parameter is omitted, the sequence value is 0, and the order of the ACEs inside the cisco-av-pair RADIUS VSA is used.

The following example is an access list definition as it should be configured for a cisco-av-pair VSA on a RADIUS server:

```
ip:inacl#1=permit tcp 10.1.0.0 255.0.0.0 10.0.0.0 255.0.0.0
ip:inacl#99=deny tcp any any
ip:inacl#2=permit udp 10.1.0.0 255.0.0.0 10.0.0.0 255.0.0.0
ip:inacl#100=deny udp any any
ip:inacl#3=permit icmp 10.1.0.0 255.0.0.0 10.0.0.0 255.0.0.0
```

For information about making unique per user the access lists that are sent in the cisco-av-pair attribute, see the documentation for your RADIUS server.

On the security appliance, the downloaded access list name has the following format:

```
AAA-user-username
```

The *username* argument is the name of the user that is being authenticated.

The downloaded access list on the security appliance consists of the following lines. Notice the order based on the numbers identified on the RADIUS server.

```
access-list  AAA-user-bcham34-79AD4A08 permit tcp 10.1.0.0 255.0.0.0 10.0.0.0 255.0.0.0
access-list  AAA-user-bcham34-79AD4A08 permit udp 10.1.0.0 255.0.0.0 10.0.0.0 255.0.0.0
access-list  AAA-user-bcham34-79AD4A08 permit icmp 10.1.0.0 255.0.0.0 10.0.0.0 255.0.0.0
access-list  AAA-user-bcham34-79AD4A08 deny tcp any any
access-list  AAA-user-bcham34-79AD4A08 deny udp any any
```

Downloaded access lists have two spaces between the word “access-list” and the name. These spaces serve to differentiate a downloaded access list from a local access list. In this example, “79AD4A08” is a hash value generated by the security appliance to help determine when access list definitions have changed on the RADIUS server.

Converting Wildcard Netmask Expressions in Downloadable Access Lists

If a RADIUS server provides downloadable access lists to Cisco VPN 3000 Series Concentrators as well as to the security appliance, you may need the security appliance to convert wildcard netmask expressions to standard netmask expressions. This is because Cisco VPN 3000 Series Concentrators support wildcard netmask expressions but the security appliance only supports standard netmask expressions. Configuring the security appliance to convert wildcard netmask expressions helps minimize the effects of these differences upon how you configure downloadable access lists on your RADIUS servers. Translation of wildcard netmask expressions means that downloadable access lists written for Cisco VPN 3000 Series Concentrators can be used by the security appliance without altering the configuration of the downloadable access lists on the RADIUS server.

You configure access list netmask conversion on a per server basis, using the **acl-netmask-convert** command, available in the aaa-server configuration mode. For more information about configuring a RADIUS server, see AAA Setup. For more information about the **acl-netmask-convert** command, see the *Cisco Security Appliance Command Reference*.

Configuring a RADIUS Server to Download Per-User Access Control List Names

To download a name for an access list that you already created on the security appliance from the RADIUS server when a user authenticates, configure the IETF RADIUS filter-id attribute (attribute number 11) as follows:

filter-id=*acl_name*

Note In Cisco Secure ACS, the value for filter-id attributes are specified in boxes in the HTML interface, omitting **filter-id=** and entering only *acl_name*.

For information about making unique per user the filter-id attribute value, see the documentation for your RADIUS server.

Configuring Filter Rules

This section contains the following topics:

URL Filtering

Configuration > Properties > URL Filtering

You can apply filtering to connection requests originating from a more secure network to a less secure network. Although you can use ACLs to prevent outbound access to specific content servers, managing usage this way is difficult because of the size and dynamic nature of the Internet. You can simplify configuration and improve security appliance performance by using a separate server running one of the following Internet filtering products:

- Websense Enterprise for filtering HTTP, HTTPS, and FTP.
- Secure Computing SmartFilter for filtering HTTP only. (Although some versions of Sentian support HTTPS, the security appliance only supports filtering HTTP with Sentian.)

Although security appliance performance is less affected when using an external server, users may notice longer access times to websites or FTP servers when the filtering server is remote from the security appliance.

When filtering is enabled and a request for content is directed through the security appliance, the request is sent to the content server and to the filtering server at the same time. If the filtering server allows the connection, the security appliance forwards the response from the content server to the originating client. If the filtering server denies the connection, the security appliance drops the response and sends a message or return code indicating that the connection was not successful.

If user authentication is enabled on the security appliance, then the security appliance also sends the user name to the filtering server. The filtering server can use user-specific filtering settings or provide enhanced reporting regarding usage.

General Procedure

The following summarizes the procedure for enabling filtering with an external filtering server.

Step 1 Identify the filtering server.

Step 2 (Optional) Buffer responses from the content server (optional).

Step 3 (Optional) Cache content server addresses to improve performance (optional).

Step 4 Configure filtering rules. See **Filter Rules**.

Step 5 Configure the external filtering server. For more information refer to the following websites:

- http://www.websense.com
- http://www.securecomputing.com

You can identify up to four filtering servers per context. In single mode a maximum of 16 servers are allowed. The security appliance uses the servers in order until a server responds. You can only configure a single type of server (Websense or Secure Computing SmartFilter) in your configuration.

Note You must add the filtering server before you can configure filtering for HTTP, HTTPS, or FTP filtering rules.

Fields

- URL Filtering Server area
  - Websense—Enables the Websense URL filtering servers
  - Secure Computing SmartFilter—Enables the Secure Computing SmartFilter URL filtering server.
  - Secure Computing SmartFilter Port—Specifies the Secure Computing SmartFilter port. The default is 4005.
  - Interface—Displays the interface connected to the filtering server.
  - IP Address—Displays the IP address of the filtering server.
  - Timeout—Displays the number of seconds after which the request to the filtering server times out.
  - Protocol—Displays the protocol used to communicate with the filtering server.
  - TCP Connections—Displays the maximum number of TCP connections allowed for communicating with the URL filtering server.
  - Add—Adds a new filtering server, depending on whether you have selected Websense or Secure Computing SmartFilter.
  - Insert Before—Adds a new filtering server in a higher priority position than the currently selected server.
  - Insert After—Adds a new filtering server in a lower priority position than the currently selected server.
  - Edit—Lets you modify parameters for the selected filtering server
  - Delete—Deletes the selected filtering server.
- Apply—Applies the changes to the running configuration.
- Reset—Removes any changes that have not been applied.
- Advanced—Displays advanced filtering parameters, including buffering caching, and long URL support.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Filter Rules

Add/Edit Parameters for Websense URL Filtering

Configuration > Properties > URL Filtering > Add/Edit Parameters for Websense URL Filtering

- Interface—Specifies the interface on which the URL filtering server is connected.
- IP Address—Specifies the IP address of the URL filtering server.
- Timeout—Specifies the number of seconds after which the request to the filtering server times out.
- Protocol area
  - TCP 1—Uses TCP Version 1 for communicating with the Websense URL filtering server.
  - TCP 4—Uses TCP Version 4 for communicating with the Websense URL filtering server.
  - UDP 4—Uses UDP Version 4 for communicating with the Websense URL filtering server.
- TCP Connections—Specifies the maximum number of TCP connections allowed for communicating with the URL filtering server.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Parameters for Secure Computing SmartFilter URL Filtering

Configuration > Properties > URL Filtering > Add/Edit Parameters for Secure Computing SmartFilter URL Filtering

- Interface—Specifies the interface on which the URL filtering server is connected.
- IP Address—Specifies the IP address of the URL filtering server.
- Timeout—Specifies the number of seconds after which the request to the filtering server times out.
- Protocol area
  - TCP—Uses TCP for communicating with the Secure Computing SmartFilter URL filtering server.

  - UDP—Uses UDP for communicating with the Secure Computing SmartFilter URL filtering server.

TCP Connections—Specifies the maximum number of TCP connections allowed for communicating with the URL filtering server.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Advanced URL Filtering

Configuration > Properties > URL Filtering > Advanced URL Filtering

Fields

URL Cache Size area

After a user accesses a site, the filtering server can allow the security appliance to cache the server address for a certain amount of time, as long as every site hosted at the address is in a category that is permitted at all times. Then, when the user accesses the server again, or if another user accesses the server, the security appliance does not need to consult the filtering server again.

Note Requests for cached IP addresses are not passed to the filtering server and are not logged. As a result, this activity does not appear in any reports.

- Enable caching based on—Enables caching based on the specified criteria.
  - Destination Address—Caches entries based on the URL destination address. Select this mode if all users share the same URL filtering policy on the Websense server.
  - Source/Destination Address—Caches entries based on both the source address initiating the URL request as well as the URL destination address. Select this mode if users do not share the same URL filtering policy on the server
  - Cache size—Specifies the size of the cache.

URL Buffer Size area

When a user issues a request to connect to a content server, the security appliance sends the request to the content server and to the filtering server at the same time. If the filtering server does not respond before the content server, the server response is dropped. This delays the web server response from the point of view of the web client because the client must reissue the request.

By enabling the HTTP response buffer, replies from web content servers are buffered and the responses are forwarded to the requesting client if the filtering server allows the connection. This prevents the delay that might otherwise occur.

- Enable buffering—Enables request buffering.
  - Number of 1550-byte buffers—Specifies the number of 1550-byte buffers.

- Long URL Support area

  By default, the security appliance considers an HTTP URL to be a long URL if it is greater than 1159 characters. For Websense servers, you can increase the maximum length allowed.

  - Use Long URL—Enables long URLs for Websense filtering servers.
  - Maximum Long URL Size—Specifies the maximum URL length allowed, up to a maximum of 4 KB.
  - Memory Allocated for Long URL—Specifies the memory allocated for long URLs.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Filter Rules

Configuration > Security Policy > Filter Rules

The Filter Rules window displays configured filter rules and provides options for adding new filter rules or modifying existing rules. A filter rule specifies the type of filtering to apply and the kind of traffic to which it should be applied.

Note Before you can add an HTTP, HTTPS, or FTP filter rule, you must enable a URL filtering server. To enable a URL filtering server, use the **Features > Configuration > Properties > URL Filtering** screen. For more information see URL Filtering.

Benefits

The Filter Rules window provides information about the filter rules that are currently configured on the security appliance. It also provides buttons that you can use to add or modify the filter rules and to increase or decrease the amount of detail shown in the window.

Filtering allows greater control over any traffic that your security policy allows to pass through the security appliance. Instead of blocking access altogether, you can remove specific undesirable objects from HTTP traffic, such as ActiveX objects or Java applets, that may pose a security threat in certain situations. You can also use URL filtering to direct specific traffic to an external filtering server, such as Secure Computing SmartFilter or Websense. These servers can block traffic to specific sites or types of sites, as specified by your security policy.

Because URL filtering is CPU-intensive, using an external filtering server ensures that the throughput of other traffic is not affected. However, depending on the speed of your network and the capacity of your URL filtering server, the time required for the initial connection may be noticeably slower for filtered traffic.

Fields

- No—Numeric identifier of the rule. Rules are applied in numeric order.

- Source—Source host or network to which the filtering action applies.
- Destination—Destination host or network to which the filtering action applies.
- Service—Identifies the protocol or service to which the filtering action applies.
- Action—Type of filtering action to apply.
- Options—Indicates the options that have been enabled for the specific action.
- Add—Displays the Add Filter Rule dialog box for adding a new filtering rule.
- Edit—Displays the Edit Filter Rule dialog box for editing the selected filtering rule.
- Delete—Deletes the selected filtering rule.
- MoveUp—Moves the filter rule up.
- MoveDown—Moves the filter rule down.
- Cut—Lets you to cut a filter rule and place it elsewhere.
- Copy—Lets you copy a filter rule.
- Paste—Lets you paste a filter rule elsewhere.
- Find—Lets you search for a filter rule. Clicking on this button brings up an extended tool bar.
  - Filter—Lets you search by source, destination, source, action, or rule query, using the drop-down menu.
  - ....—Lets you select the source of the filter, and brings up the Select Source dialog box.
  - Filter—Lets you input a filter.
  - Clear—Lets you clear a filter rule.
  - Rule Query—Lets you devise a query to search for a rule.
- Use the Addresses tab to select the source of the filter rule that you are choosing.
  - Type—Lets you select a source from the drop-down menu, selecting from All, IP Address Objects, IP Names, or Network Object groups.
  - Name—Lists the name(s) of the filter rule.
  - Add—Lets you add a filter rule.
  - Edit—Lets you edit a filter rule.
  - Delete—Lets you delete a filter rule.
  - Find—Lets you find a filter rule.
- Use the Services tab to select a predefined filter rule.
  - Type—Lets you select a source from the drop-down menu, selecting from All, IP Address Objects, IP Names, or Network Object groups.
  - Name—Lists the name(s) of the filter rule.
  - Edit—Lets you edit a filter rule.
  - Delete—Lets you delete a filter rule.
  - Find—Lets you find a filter rule.
- Use the Time Ranges to select a time range for the filter rule.
  - Add—Add—Lets you add a time range for the filter rule.
  - Edit—Lets you edit a time range for the filter rule.

– Delete—Lets you delete a time range for a filter rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Select Source

Configuration > Security Policy > Filter Rules >Select Source

Use the Select Source dialog box to select the source of the filter rule that you are closing.

Fields

- Type—Lets you select a source from the drop-down menu, selecting from All, IP Address Objects, IP Names, or Network Object groups.
- Name—Lists the name(s) of the filter rule.
- IP Address—Lists the IP address of the filter rule(s).
- Netmask—Lists the netmask of the filter rule(s).
- Description (optional)—Lists descriptions for the filter rules.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Rule Query

Configuration > Security Policy > Filter Rules >Select Source > Rule Query

Fields

- Name—Lets you enter the name of the filter rule for the query.
- Description (optional)—Lets you enter a description of the filter rule for the query.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Filter Rule

Configuration > Security Policy > Filter Rules > Add/Edit Filter Rule

Use the Add Filter Rule dialog box to specify the interface on which the rule applies, to identify the traffic to which it applies, or to configure a specific type of filtering action.

Note Before you can add an HTTP, HTTPS, or FTP filter rule, you must enable a URL filtering server. To enable a URL filtering server, use the **Features > Configuration > Properties > URL Filtering** screen. For more information see URL Filtering.

Fields

- Action—Provides the following drop-down list of different filtering actions to apply:
  - Filter ActiveX
  - Do not filter ActiveX
  - Filter Java Applet
  - Do not filter Java Applet
  - Filter HTTP (URL)
  - Do not filter HTTP (URL)
  - Filter HTTPS
  - Do not filter HTTPS
  - Filter FTP
  - Do not filter FTP

  The Rule Flow Diagram and the Filtering Option area changes according to which filtering action you select.
- Source area
  - IP Address—Use the IP address to identify the traffic to which the filtering action applies.
  - ...—Opens the Browse Source Address dialog box.
  - Netmask—Specifies the Subnet mask used to identify the traffic to which the filtering action applies when IP Address is selected.
- Destination area
  - IP Address—Identifies the traffic to which the filtering action applies.
  - Netmask—Specifies the Subnet mask used to identify the traffic to which the filtering action applies when IP Address is selected.

- Rule Flow Diagram area —Provides a graphic representation of how a specific filtering action is applied to traffic that is forwarded through the security appliance.
- ActiveX Filtering Option area—This area appears only when you select the Filter ActiveX option from the drop-down list.
  - ActiveX Filtering Option—When you select the Filter ActiveX option from the drop-down list, this field appears and lets you specify the TCP/UDP port on which the security appliance listens for traffic to which the filtering action applies.
- Java Filtering Option—This area appears only when you select the Filter Java option from the drop-down list.
  - Java Filtering Option—When you select the Filter Java option from the drop-down list, this field appears and lets you specify the TCP/UDP port on which the security appliance listens for traffic to which the filtering action applies.
- HTTP Filtering Option—This area appears only when you select the Filter HTTP option from the drop-down list.
  - Filter HTTP on port(s)—Specify the TCP/UDP port on which the security appliance listens for traffic to which the filtering action applies.
  - Block connections to proxy server—Prevent HTTP requests made through a proxy server.
  - Allow outbound traffic if URL server is not available—When enabled, if the URL filtering server is down or connectivity is interrupted to the security appliance, users will be able to connect without URL filtering being performed. If this is disabled, users will not be able to connect to Internet websites when the URL server is unavailable.
  - Truncate CGI requests by removing the CGI parameters—The security appliance forwards only the CGI script location and the script name, without any parameters, to the filtering server.
- HTTPS Filtering Option—This area appears only when you select the Filter HTTPS option from the drop-down list.
  - Filter HTTPS on port(s)—specify the TCP/UDP port on which the security appliance listens for traffic to which the filtering action applies.
  - Allow outbound traffic if URL server is not available—When enabled, if the URL filtering server is down or connectivity is interrupted to the security appliance, users will be able to connect without URL filtering being performed. If this is disabled, users will not be able to connect to Internet websites when the URL server is unavailable.
- FTP Filtering Option—This area appears only when you select the Filter FTP option from the drop-down list.
  - Filter FTP on port(s)—Specifies the TCP/UDP port on which the security appliance listens for traffic to which the filtering action applies.
  - Allow outbound traffic if URL server is not available—When enabled, if the URL filtering server is down or connectivity is interrupted to the security appliance, users will be able to connect without URL filtering being performed. If this is disabled, users will not be able to connect to Internet websites when the URL server is unavailable.
  - Block outbound traffic if absolute FTP path is not provided—When enabled, FTP requests are dropped if they use a relative path name to the FTP directory.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Browse Source/Destination Address

Configuration > Security Policy > Filter Rules > Add/Edit Filter Rule >Browse Source Address

Fields

- Type—Lets you select from one of the following types of sources: IP Address Objects, IP Names, or Network Address Groups.
- Name—Specifies the name used to identify the traffic to which the filtering action applies when the Name button is selected.
- IP Address—Specifies the IP address used to identify the traffic to which the filtering action applies.
- Netmask—Specifies the Subnet mask used to identify the traffic to which the filtering action applies when IP Address is selected.
- Description (optional)—Specifies a description for the filter.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Filter Rules

URL Filtering

Configuring Service Policy Rules

This chapter describes how to enable service policy rules. Service policy rules define how specific types of application inspection are applied to different types of traffic that is received by the security appliance. You apply a specific rule to an interface or globally to every interface.

Configuring Service Policy Rules

This section describes how to configure service policy rules, and includes the following topics:

Service Policy Rules

Configuration > Security Policy > Service Policy Rules

Some applications require special handling by the security appliance and specific application inspection engines are provided for this purpose. Applications that require special application inspection engines are those that embed IP addressing information in the user data packet or open secondary channels on dynamically assigned ports. Application inspection is enabled by default for many protocols, while it is disabled for other protocols. In many cases, you can change the port on which the application inspection listens for traffic.

Application inspection engines work with NAT to help identify the location of embedded addressing information. This allows NAT to translate these embedded addresses and to update any checksum or other fields that are affected by the translation.

Service policy rules define how specific types of application inspection are applied to different types of traffic that is received by the security appliance. You apply a specific rule to an interface or globally to every interface.

Use traffic match criteria to define the set of traffic to which you want to apply application inspection. For example, TCP traffic with a port value of 23 might be classified as the Telnet traffic class. You can use the traffic class to change the default port for application inspection for protocols where this is permitted.

Multiple traffic match criteria can be assigned to a single interface, but a packet will only match the first criteria within a specific service policy rule.

Fields

- Add—Adds a new service policy rule. Choose the type of rule you want to add from the drop-down list.
- Edit—Edits a service policy rule.
- Delete—Deletes a service policy rule.
- Move Up—Moves a rule up. Rules are assessed in the order they appear in this table, so the order can matter if you have overlapping rules.
- Move Down—Moves a rule down.
- Cut—Cuts a rule.
- Copy—Copies the parameters of a rule so you can start a new rule with the same parameters using the Paste button.
- Paste—Opens an Add/Edit Rule dialog box with the copied or cut parameters of a rule prefilled. You can then make any modifications and add it to the table. The Paste button adds the rule above the selected rule. The Paste After item, available from the Paste drop-down list, adds the rule after the selected rule.
- Find—Filters the display to show only matching rules. Clicking **Find** opens the Filter field. Click **Find** again to hide the Filter field.
  - Filter drop-down list—Choose the criteria to filter on, either Interface, Source, Destination, Service, or Rule Query. A rule query is a collection of multiple criteria that you can save and use repeatedly.
  - Filter field—For the Interface type, this field becomes a drop-down list so you can choose an interface name, or **All Interfaces**. For the Rule Query type, the drop-down list includes all defined rule queries. The Source and Destination types accept an IP address. You can type one manually, or browse for one by clicking the **...** button and launching the Browse Address dialog box. The Service type accepts a TCP, UDP, TCP-UDP, ICMP, or IP protocol type. You can type one manually, or browse for one by clicking the **...** button and launching the Browse Service Groups dialog box.
  - Filter—Runs the filter.
  - Clear—Clears the Filter field.
  - Rule Query—Opens the Rule Queries dialog box so you can manage named rule queries.
- Show Rule Flow Diagram—Shows the Rule Flow Diagram area under the rule table. This diagram shows the networks, type of traffic, interface name, direction of flow, and action.
- Packet Trace—Opens the Packet Tracer tool with the parameters pre-filled with the characteristics of the selected rule.

The following description summarizes the columns in the Service Policy Rules table. You can edit the contents of these columns by double-clicking on a table cell. Double-clicking on a column header sorts the table in ascending alphanumeric order, using the selected column as the sort key. If you right-click a rule, you see all of the options represented by the buttons above, as well as Insert and Insert After items. These items either insert a new rule before the selected rule (Insert) or after the selected rule (Insert After.)

- Name—Indicates the name of the rule.
- No—Indicates the order of evaluation for the rule.
- Enabled—Indicates whether the rule is enabled or disabled.
- Match—Indicates if the criteria are used to include (match) or exclude (do not match) traffic.

- Source—Lists the IP addresses that are subject to service policy when traffic is sent to the IP addresses listed in the Destination column.
- Destination—Lists the IP addresses that are subject to service policy when traffic is sent from the IP addresses listed in the Source column.
- Service—Shows the service or protocol specified by the rule.
- Time—Displays the time range during which the rule is applied.
- Rule Actions—Shows the actions applied by the rule.
- Description—The description you entered when you added the rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Service Policy

Add Service Policy Rule Wizard > Service Policy

The Service Policy dialog box lets you add a new service policy rule, apply the rule to a specific interface, or apply the rule globally to all interfaces.

Fields

- Create a service policy and apply to area
  - Interface—Applies the rule to a specific interface. This selection is required if you want to match traffic based on the source or destination IP address using an access list.
  - Interface—Specifies the interface to which the rule applies.
  - Policy Name—Specifies the name of the interface service policy.
  - Description—Provides a text description of the policy.
  - Global - applies to all interfaces—Applies the rule to all interfaces. This selection is not compatible with matching traffic based on the source or destination IP address using an access list.
  - Policy Name—Specifies the name of the global service policy. Only one global service policy is allowed and it cannot be renamed.
  - Description—Provides a text description of the policy.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Edit Service Policy

Configuration > Security Policy > Service Policy Rules > Edit Service Policy

The Edit Service Policy dialog box lets you change the description for the selected service policy.

Fields

- Description—Provides a text description of the service policy.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Traffic Classification Criteria

Add/Edit Service Policy Rule Wizard > Traffic Classification Criteria

The Traffic Classification tab on the Edit Service Policy Rule screen lets you specify the criteria you want to use to match traffic to which the security policy rule applies.

Fields

- Name—Identifies the name of the traffic class.
- Description (optional)—Provides a text description of the new traffic class.
- Traffic match criteria area:
  - Default Inspection Traffic—Uses the criteria specified in the default inspection traffic policy.
  - Source and Destination IP Address (uses ACL)—Matches traffic based on the source and destination IP address, using an ACL. This selection is only available if you apply the rule to a specific interface using an Interface Service Policy.
  - Tunnel Group—Matches traffic based on the tunnel group.
  - TCP or UDP Destination Port—Matches traffic based on the TCP or UDP destination port.
  - RTP Range—Matches traffic based on a range of RTP ports.
  - IP DiffServ CodePoints (DSCP)—Matches traffic based on the Differentiated Services model of QoS.

  - IP Precedence—Matches traffic based on the IP precedence model of QoS.
  - Any traffic—Matches all traffic regardless of the traffic type.
- Add rule to existing traffic class—Adds the rule to the existing traffic class that is selected in the drop-down list.
- Use class-default as the traffic class—Specifies that the class-default traffic class is used when traffic does not match any other traffic class.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Default Inspections

Add/Edit Service Policy Rule Wizard > Traffic Classification Criteria > Default Inspections

The Default Inspections dialog box lists the default port assignments that are used when you select the Default Inspection Traffic criteria on the Traffic Classification Criteria dialog box.

- Service—This lists the application inspection engine type.
- Protocol—This identifies whether the application inspection uses TCP or UDP for the transport protocol.
- Port—This identifies the port number used by the Default Inspection Traffic criteria.

Management Type Traffic Class and Action

Add/Edit Service Policy Rule Wizard > Management Type Traffic Class and Action

The Management Class dialog box lets you configure the management traffic classification and define actions for the classified traffic.

Fields

- Name—Identifies the name of the traffic management class.
- Description (optional)—Provides a text description of the new traffic management class.
- Match on Destination Port area:
  - Protocol—Matches traffic based on the TCP or UDP destination port.
  - Service—Choose the = (equal) operator or range to specify a range of ports. Either type a number, or choose a well-known port name from the drop-down list. For a range, you must specify numbers.
- Protocol Inspection area:
  - RADIUS Accounting Map—Choose a defined RADIUS accounting map from the drop-down list.

- Configure—Opens the Select RADIUS Accounting Map dialog box to select a defined RADIUS accounting map or add a RADIUS accounting map or for fine control over inspection.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Select RADIUS Accounting Map

Add/Edit Service Policy Rule Wizard > Management Type Traffic Class and Action > 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add RADIUS Accounting Policy Map

Add/Edit Service Policy Rule Wizard > Management Type Traffic Class and Action > Select 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).

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Using Default Inspection Traffic Criteria

The **fixup** command, which is available in PIX Firewall Version 6.3 and earlier releases, provided a simple, global policy for application inspection. The Modular Policy Framework provides a much more granular method of inspecting traffic. Modular Policy Framework lets you select the traffic for a specific application inspection and this can improve the performance of the security appliance. Performance is improved because the application inspection engine only inspects a limited amount of traffic.

To simplify enabling application inspection on the default ports, use the default inspection traffic criteria. When you specify the default inspection traffic criteria the security appliance selects traffic for application inspection on the well-known port for each protocol. Table 1 lists the default port assignments for each protocol.

Table 1 Default Port Assignments

Protocol Name	Protocol	Source Port	Destination Port
ctiqbe	tcp	N/A	2748
dns	udp	53	53
esmtp/smtp	tcp	N/A	25
ftp	tcp	N/A	21
gtp	udp	2123,3386	2123,3386
h323 h225	tcp	N/A	1720
h323 ras	udp	N/A	1718-1719
http	tcp	N/A	80
icmp	icmp	N/A	N/A

Table 1 Default Port Assignments (continued) (continued)

Protocol Name	Protocol	Source Port	Destination Port
ils	tcp	N/A	389
mgcp	udp	2427,2727	2427,2727
netbios	udp	137-138	N/A
pptp	tcp	1723	1723
rsh	tcp	N/A	514
rtsp	tcp	N/A	554
sip	tcp, udp	N/A	5060
skinny	tcp	N/A	2000
sqlnet	tcp	N/A	1521
sunrpc	udp	111	111
tftp	udp	N/A	69
xdmcp	udp	177	177

When you select the default inspection traffic criteria, you can then enable each protocol on the **Protocol Inspection** tab of the **Rule Actions** screen. The protocol will be enabled on its default port. You can restrict inspection to a specific flow by using the **Source and destination IP address (uses ACL)** button and selecting specific criteria, such as **Source Host/Network** or **Destination Host/Network** from the **Service Policy Rule** screen.

Note The default inspection traffic criteria override any port settings in the Protocol and Service group box. That means that you cannot change any of the default port assignments for any protocol when the default inspection traffic criteria are used.

The inspection_default security policy is a preconfigured global policy that enables application inspection using the default inspection traffic criteria. This global policy is enabled in the security appliance factory default configuration.

Note When you specify the default inspection traffic as the traffic match criteria, only inspect rule actions can be applied in the security policy for the specified interface. Actions on the QoS and Connection Settings tabs cannot be applied.

Changing Default Ports for Application Inspection

The default inspection traffic criteria override any port settings in the Protocol and Service group box. That means that you cannot change any of the default port assignments for any protocol when the default inspection traffic criteria are used.

To change the default port assignment for any protocol, you must manually configure and enable each inspection engine.

To use Modular Policy Framework for changing the default port assignment for a protocol, perform the following steps:

Step 1 Click **Service Policy Rules** on the **Security Policy** panel and then click **Add**.

The **Add Service Policy Rule Wizard - Service Policy** screen appears.

Step 2 Create a service policy.

To create a security policy for a specific interface, on the **Create a service policy and apply to** group box, click the **Interface** radio button and select an available interface from the selection list.

To create a global security policy to be applied to all interfaces, on the **Create a service policy and apply to** group box, click the **Global** radio button.

Step 3 Type a name of up to 40 characters in the **Policy Name** box and click **Next**.

The **Add Service Policy Rule Wizard - Traffic Classification Criteria** screen appears.

Step 4 Click the **Source and destination IP address (uses ACL)** button.

Step 5 Select the **Source Port** and **Destination Port** for the protocol in the **Protocol and Service** group box and click **Next**.

The **Add Service Policy Rule Wizard - Rule Actions** screen appears.

Step 6 Click the checkbox for the protocol you want to enable and click **Finish**.

The new service policy is shown in the **Service Policy Rules** table on the **Security Policy** panel.

Step 7 To enable another inspection engine, select the service policy and click **Add**.

The **Add Service Policy Rule Wizard - Service Policy** screen appears.

Step 8 Click **Next**.

The **Add Service Policy Rule Wizard - Traffic Classification Criteria** screen appears.

Step 9 Click **Create a new traffic class** and change the name of the traffic class, if necessary.

By default, the number at the end of the name for each traffic class is incremented as you add each new class.

Step 10 Click **Source and destination IP address (uses ACL)**.

Step 11 Click the **Traffic Match** tab.

Step 12 Select the second port number for the protocol in the **Protocol and Service** group box and click **OK**.

The new access control entry is shown in the **Service Policy Rules** table on the **Security Policy** panel.

Configuring Application Inspection with Multiple Ports

To use Modular Policy Framework for changing the default port assignment for protocols that use more than one port, perform the following steps:

Step 1 Click **Service Policy Rules** on the **Security Policy** panel and then click **Add**.

The **Add Service Policy Rule Wizard - Service Policy** screen appears.

Step 2 Create a service policy.

To create a security policy for a specific interface, on the **Create a service policy and apply to** group box, click the **Interface** radio button and select an available interface from the selection list.

To create a global security policy to be applied to all interfaces, on the **Create a service policy and apply to** group box, click the **Global** radio button.

Step 3 Type a name of up to 40 characters in the **Policy Name** box and click **Next**.

The **Add Service Policy Rule Wizard - Traffic Classification Criteria** screen appears.

Step 4 Click the **Source and destination IP address (uses ACL)** button.

Step 5 Select the first port number for the protocol in the **Protocol and Service** group box and click **Next**.

The **Add Service Policy Rule Wizard - Rule Actions** screen appears.

Step 6 Define the rule action to apply to the specified traffic flow, using one of the following tabs:

- **Protocol Inspection**
- **Connection Settings**
- **QoS**

Step 7 Click **Finish**.

The new service policy is shown in the **Service Policy Rules** table on the **Security Policy** panel.

Step 8 Right-click the security policy on the Service Policy Rules table.

Step 9 On the pop-up menu that appears, select **Insert After**.

The **Insert Service Policy Rule After** screen appears.

Step 10 Click the **Traffic Match** tab.

Step 11 Select the second port number for the protocol in the **Protocol and Service** group box and click **OK**.

The new access control entry is shown in the **Service Policy Rules** table on the **Security Policy** panel.

Source and Destination Address (This dialog is called "ACL" in other contexts)

(This dialog box is called ACL when editing a service policy rule)

- **Add/Edit Service Policy Rule Wizard > Traffic Match > Source and Destination Address**
- **Configuration > Security Policy > Edit Service Policy Rule > ACL Tab**

(You can get to this dialog box through various paths.)

This dialog box lets you identify the traffic to which a service policy rule applies based on the IP address or TCP/UDP port of the sending or receiving host. You can also use this dialog box to select a Time Range during which the policy rule is in effect.

Fields

- Select an action—Lets you specify whether the traffic must match or must not match the criteria specified on this dialog box.
- Time Range area
  - Time Range—Lets you select a named time range during which the policy rule is in effect.
  - New—Lets you access the Add Time Range dialog box. For more information, see Add/Edit Time Range.
- Source Host/Network area

  - IP Address—Specifies that the source of the traffic is to be identified by IP address. When you select this button, the Interface drop-down list, IP address field, . . . button, and Mask drop-down list appear within the area.
  - Name—Specifies that the source of the traffic is to be identified by interface name. When you select this button, the Name drop-down list appears within the area.
  - Group—Specifies that the source of the traffic is to be identified by object groups. When you select this button, the Interface drop-down list and Group drop-down list appear within the area.
  - Interface—Specifies the name of the interface that the source of the traffic is on. This drop-down list appears only when the IP Address button or the Group button is selected.
  - IP address—Specifies the IP address used to identify the source of the traffic. This field appears only when the IP Address button is selected.
  - . . .—Lets you access the Select host/network dialog box, which lets you select a host or network from a preconfigured drop-down list. This button appears only when the IP Address button is selected.
  - Mask—Specifies the subnet mask for the address entered in the IP address field. This field appears only when the IP Address button is selected.
  - Name—Specifies the name of the interface that the source of the traffic is on. This drop-down list appears only when the Name button is selected.
  - Group—Specifies the object group that the source of the traffic is in. The items on the drop-down list is controlled by the Hosts/Networks window. For more information about that window, see Network Object Groups. The Group drop-down list appears only when the Group button is selected.

- Destination Host/Network area
  - IP Address—Specifies that the destination of the traffic is to be identified by IP address. When you select this button, the Interface drop-down list, IP address field, . . . button, and Mask drop-down list appear within the area.
  - Name—Specifies that the destination of the traffic is to be identified by interface name. When you select this button, the Name drop-down list appears within the area.
  - Group—Specifies that the destination of the traffic is to be identified by object groups. When you select this button, the Interface drop-down list and Group drop-down list appear within the area.
  - Interface—Specifies the name of the interface that the destination of the traffic is on This drop-down list appears only when the IP Address button or the Group button is selected.
  - IP address—Specifies the IP address used to identify the destination of the traffic. This field appears only when the IP Address button is selected.
  - . . .—Lets you access the Select host/network dialog box, which lets you select a host or network from a preconfigured drop-down list. This button appears only when the IP Address button is selected.
  - Mask—Specifies the subnet mask for the address entered in the IP address field. This field appears only when the IP Address button is selected.
  - Name—Specifies the name of the interface that the destination of the traffic is on. This drop-down list appears only when the Name button is selected.

  - Group—Specifies the object group that the destination of the traffic is in. The items on the drop-down list is controlled by the Hosts/Networks window. For more information about that window, see Network Object Groups. The Group drop-down list appears only when the Group button is selected.

- Rule Flow Diagram—Provides a graphic representation of how a specific filtering action is applied to traffic that is forwarded through the security appliance.
- Protocol and Service area
  - TCP—Matches traffic based on the TCP protocol or service.
  - UDP—Matches traffic based on the UDP protocol or service.
  - ICMP—Matches traffic based on the ICMP protocol value.
  - IP—Matches traffic based on the IP protocol value.
  - Manage Service Groups—Displays the Manage Service Groups dialog box, which lets you create and edit service groups. This button is available only when the TCP button is selected.
  - Source Port—Appears only when either the TCP or UDP radio button is selected.

    Service—Matches traffic based on the source port value.

    Operator—Specifies whether to identify a singe port or a range of ports to match. When you select = (equal to), not= (not equal to), > (greater than), or < (less than) from the drop-down list, the . . . button appears, which lets you select a specific named port. When you select range from the drop-down list, two fields appear that let you enter the starting and ending ports in the range.

    ...—Displays the Service dialog box, which lets you select the named values for the TCP or UDP ports to match.

    Service Group—Matches traffic based on the source service group. To control the items on the drop-down list, use the Manage Service Groups button.
  - Destination Port—Appears only when either the TCP or UDP radio button is selected.

    Service—Matches traffic based on the destination port value.

    Operator—Specifies whether to identify a singe port or a range of ports to match. When you select = (equal to), not= (not equal to), > (greater than), or < (less than) from the drop-down list, the . . . button appears, which lets you select a specific named port. When you select range from the drop-down list, two fields appear that let you enter the starting and ending ports in the range.

    ...—Displays the Service dialog box, which lets you select the named values for the TCP or UDP ports to match.

    Service Group—Matches traffic based on the destination service group. To control the items on the drop-down list, use the Manage Service Groups button.
  - ICMP Type—Appears only when the ICMP radio button is selected.

    ICMP type—Lets you enter the ICMP type of the traffic.

    ...—Displays the Service dialog box, which lets you select ICMP types from a preconfigured drop-down list.
  - IP Protocol—Appears only when the IP radio button is selected.

    IP protocol—Lets you enter the IP protocol of the traffic.

    ...—Displays the Service dialog box, which lets you select an IP protocol from a preconfigured drop-down list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Destination Port

Add/Edit Service Policy Rule Wizard > Traffic Match > Destination Port

The Destination Port dialog box appears when you select TCP or UDP destination port in the Traffic Match Criteria dialog box, or choose the corresponding tab when editing a service policy rule. This dialog box lets you identify the traffic to which a service policy rule applies based on the TCP or UDP destination port.

Fields

- TCP—Matches traffic based on the TCP port used by the destination.
- UDP—Matches traffic based on the UDP port used by the destination.
- Operator—Specifies whether to identify a singe port or a range of ports to match.

  When you select = (equals sign) from the drop-down list, the . . . button appears, which lets you select a specific named port.

  When you select range from the drop-down list, two fields appear that let you enter the starting and ending ports in the range.
- ...—Displays the Service dialog box, which lets you select the named values for the TCP or UDP ports to match.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

RTP Ports

Add/Edit Service Policy Rule Wizard > Traffic Match > RTP Ports

The RTP Ports dialog box appears when you select RTP range on the Traffic Match Criteria dialog box, or choose the corresponding tab when editing a service policy rule. This dialog box lets you identify the traffic to which a service policy rule applies based on a range of RTP ports.

- RTP Port Range—Specifies the starting and ending ports within the range of RTP ports to be used for matching traffic. RTP port numbers should be between 2000 and 65535. The maximum number of RTP ports in a range is 16383.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

IP Precedence

Add/Edit Service Policy Rule Wizard > Traffic Match > IP Precedence

The IP Precedence dialog box appears when you select IP Precedence on the Traffic Match Criteria dialog box, or choose the corresponding tab when editing a service policy rule. This dialog box lets you identify the traffic to which a service policy rule applies based on the IP precedence.

Fields

- Available IP Precedence—Lists the available IP Precedence values that you can use to match traffic. IP Precedence is one model for assigning QoS priorities to IP traffic.
- Add—Adds the selected IP Precedence value to the Match on IP Precedence list.
- Delete—Removes the selected IP Precedence value from the Match on IP Precedence list.
- Match On IP Precedence—Lists the IP Precedence values that have been selected to match traffic.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

IP DiffServ CodePoints (DSCP)

Add/Edit Service Policy Rule Wizard > Traffic Match > IP DiffServ CodePoints (DSCP)

The IP DiffServ Code Points (DSCP) dialog box lets you match traffic based on the values assigned for Differentiated Services model of QoS. DiffServ defines two sets of DSCP values: EF and AF.

Fields

- Expedited Forwarding (EF)—Provides a single DSCP value (101110) that gives marked packets the highest level of service from the network. EF is commonly considered most appropriate for Voice over IP (VoIP).
- Assured Forwarding (AF)—Provides four classes, each with three drop precedence levels.

You can select named DSCP values from the selection drop-down list, or enter a numeric value.

- Named DSCP Values—Lists named DSCP values that you can select as match criteria. Select the values that you want to match and choose **Add**.
- Enter DSCP value (0-63)—Specifies a numeric DSCP value.
- Add—Adds a selected DSCP value to the Match on DSCP table.
- Delete—Removes a selected DSCP value from the Match on DSCP table.
- Match on DSCP—Lists the DSCP values that have been selected as match criteria.
- Enter DSCP value (0-63)—Uses a numeric value as the criteria for matching.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Rule Actions > Protocol Inspection Tab

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab (You can get to this tab through various paths.)

The Protocol Inspection tab lets you enable or disable the different types of application inspection available. To view or change the configuration for a specific application inspection type, choose **Configure**, which lets you select a map name to use for the protocol. To configure a map see Configuring Inspect Maps, page 6-30.

Fields

- CTIQBE—Enables application inspection for the CTIQBE protocol.
- DCERPC—Enables application inspection for the DCERPC protocol.
  - Configure—Displays the Select DCERPC Map dialog box, which lets you select a map name to use for this protocol.
- DNS—Enables application inspection for the DNS protocol.
  - Configure—Displays the Select DNS Map dialog box, which lets you select a map name to use for this protocol.
- ESMTP—Enables application inspection for the ESMTP protocol
  - Configure—Displays the Select ESMTP Map dialog box, which lets you select a map name to use for this protocol.
- FTP—Enables application inspection for the FTP protocol.
  - Configure—Displays the Select FTP Map dialog box, which lets you select a map name to use for this protocol.
- GTP—Enables application inspection for the GTP protocol.
  - Configure—Displays the Select GTP Map dialog box, which lets you select a map name to use for this protocol.

Note GTP inspection is not available without a special license.

- H323 H225—Enables application inspection for the H323 H225 protocol.
  - Configure—Displays the Select H323 H225 Map dialog box, which lets you select a map name to use for this protocol.
- H323 RAS—Enables application inspection for the H323 RAS protocol.
  - Configure—Displays the Select H323 RAS Map dialog box, which lets you select a map name to use for this protocol.
- HTTP—Enables application inspection for the HTTP protocol.
  - Configure—Displays the Select HTTP Map dialog box, which lets you select a map name to use for this protocol.
- ICMP—Enables application inspection for the ICMP protocol.
- ICMP Error—Enables application inspection for the ICMP Error protocol.
- ILS—Enables application inspection for the ILS protocol.
- IM—Enables application inspection for the IM protocol.
  - Configure—Displays the Select IM Map dialog box, which lets you select a map name to use for this protocol.
- IPSec-Pass-Thru—Enables application inspection for the IPSec protocol.
  - Configure—Displays the Select IPSec Map dialog box, which lets you select a map name to use for this protocol.
- MGCP—Enables application inspection for the MGCP protocol.
  - Configure—Displays the Select MGCP Map dialog box, which lets you select a map name to use for this protocol.
- NETBIOS—Enables application inspection for the NetBIOS protocol.
  - Configure—Displays the Select NETBIOS Map dialog box, which lets you select a map name to use for this protocol.
- PPTP—Enables application inspection for the PPTP protocol.
- RSH—Enables application inspection for the RSH protocol.
- RTSP—Enables application inspection for the RTSP protocol.
- SCCP SKINNY—Enables application inspection for the Skinny protocol.
  - Configure—Displays the Select SCCP (Skinny) Map dialog box, which lets you select a map name to use for this protocol.
- SIP—Enables application inspection for the SIP protocol.
  - Configure—Displays the Select SIP Map dialog box, which lets you select a map name to use for this protocol.
- SNMP—Enables application inspection for the SNMP protocol.
  - Configure—Displays the Select SNMP Map dialog box, which lets you select a map name to use for this protocol.
- SQLNET—Enables application inspection for the SQLNET protocol.
- SUNRPC—Enables application inspection for the SunRPC protocol.

- TFTP—Enables application inspection for the TFTP protocol.
- XDMCP—Enables application inspection for the XDMCP protocol.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Configuring Inspect Maps

Inspect command pages for each protocol in the *Cisco Security Appliance Command Reference*

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

- No DCERPC map for inspection—Specifies no 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Configure DNS

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab > Configure DNS

Fields

Maximum DNS packet length (default 512)—Changes the maximum packet length for DNS messages that are allowed to pass through the security appliance.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Select DNS Map

Add/Edit Service Policy Rule Wizard > Rule Actions > Protocol Inspection Tab >Select DNS Map

The Select DNS Map dialog box lets you select or create a new DNS map. A DNS map lets you change the configuration values used for DNS application inspection. The Select DNS Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- No DNS map for inspection—Specifies no DNS map.
- Select a DNS 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Select ESMTP Map

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

- No ESMTP map for inspection—Specifies no 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Select FTP Map

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 security appliance to drop the connection when an embedded command is included in an FTP request.
- No FTP map for inspection—Specifies no 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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 security appliance without the required license, the security appliance displays an error message.

Fields

- No GTP map for inspection—Specifies no 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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

- No H.323 map for inspection—Specifies no 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Select HTTP Map

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

- No HTTP map for inspection—Specifies no 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Select IM Map

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

- No IM map for inspection—Specifies no IM map.
- Select an IM 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Select IPSec-Pass-Thru Map

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

- No IPSec map for inspection—Specifies no 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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

- No MGCP map for inspection—Specifies no 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Select NETBIOS Map

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

- No IM map for inspection—Specifies no 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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

- No SCCP (Skinny) map for inspection—Specifies no 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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

- No SIP map for inspection—Specifies no 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

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

- No SNMP map for inspection—Specifies no 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Rule Actions > Intrusion Prevention Tab

Add/Edit Service Policy Rule Wizard > Rule Actions > Intrusion Prevention Tab

The Intrusion Prevention tab lets you configure the Intrusion Prevention (IPS) action to take within a policy map for a traffic class. This window appears only if Intrusion Prevention System hardware is installed in the security appliance.

Fields

- Enable IPS for this traffic flow—Enables or disables intrusion prevention for this traffic flow. When this check box is selected, the other parameters on this window become active.
- Mode—Configures the operating mode for intrusion prevention
  - Inline Mode—Selects Inline Mode, in which a packet is directed to IPS. The packet might be dropped as a result of the IPS operation.
  - Promiscuous Mode—Selects Promiscuous Mode, in which IPS operates on a duplicate of the original packet. The original packet cannot be dropped.
- If IPS card fails, then—Configures the action to take if the IPS card becomes inoperable.
  - Permit traffic—Permit traffic if the IPS card fails
  - Close traffic—Block traffic if the IPS card fails.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Rule Actions > CSC Scan Tab

Add/Edit Service Policy Rule Wizard > Rule Actions > CSC Scan Tab

The CSC Scan tab lets you whether the Content Security and Control (CSC) SSM scans traffic identify by the current traffic class. This window appears only if a CSC SSM is installed in the security appliance.

The CSC SSM scans only HTTP, SMTP, POP3, and FTP traffic. If your service policy selects traffic that includes other protocols in addition to these four, packets for other protocols are passed through the CSC SSM without being scanned.

To reduce the load on the CSC SSM, configure the service policy rules that send packets to the CSC SSM to select only HTTP, SMTP, POP3, or FTP packets.

Fields

- Enable CSC scan for this traffic flow—Enables or disables use of the CSC SSM for this traffic flow. When this check box is selected, the other parameters on this window become active.
- If CSC card fails, then—Configures the action to take if the CSC SSM becomes inoperable.
  - Permit traffic—Permit traffic if the CSC SSM fails
  - Close traffic—Block traffic if the CSC SSM fails.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Managing the CSC SSM

Rule Actions > Connection Settings Tab

Add/Edit Service Policy Rule Wizard > Rule Actions > Connection Settings Tab (You can get to this tab through various paths.)

The Connection Settings tab lets you configure maximum connections, embryonic connections, and sequence number randomizing for TCP packets on a host or network. You can also configure connection timeouts and TCP normalization.

Fields

- 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 65,536. 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 65,536. 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. When a new connection is attempted by a client that already has opened the maximum per-client number of connections, the security appliance rejects the connection and drops the packet.
  - Per Client Embryonic Connections—Specifies the maximum number of simultaneous TCP embryonic connections for each client. When a new TCP connection is requested by a client that already has the maximum per-client number of embryonic connections open through the security appliance, the security appliance proxies the request to the TCP Intercept feature, which prevents the connection.
- Randomize Sequence Number—Sets the state of the Randomize Sequence Number feature to enabled or disabled. Disable this feature only if another inline security appliance is also randomizing sequence numbers and the result is scrambling the data. Each TCP connection has two Initial Sequence Numbers: one generated by the client and one generated by the server. The security appliance randomizes the ISN that is generated by the host/server on the higher security interface. At least one of the ISNs must be randomly generated so that attackers cannot predict the next ISN and potentially hijack the session.

- TCP Timeout area
  - Connection Timeout—Specifies 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.
  - Send reset to TCP endpoints before timeout—Specifies that the security appliance 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 connection slot is freed. Enter 0:0:0 to disable timeout for the connection. The default is 30 seconds.
  - Half Closed Connection Timeout—Specifies the idle time until a half closed 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 10 minutes.
- TCP Normalization area
  - Use TCP Map—Selects whether TCP normalization is enabled or not. Enable this feature to use TCP maps.
  - TCP Map—Selects an existing TCP map.
  - New—Adds a new TCP map.
  - Edit—Edits an existing TCP map.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Rule Actions > QoS Tab

Add/Edit Service Policy Rule Wizard > Rule Actions > QoS Tab (You can get to this tab through various paths.)

The QoS tab lets you apply strict scheduling priority and rate-limit traffic.

Restrictions

If a service policy is applied or removed from an interface that has existing VPN client/LAN-to-LAN or non-tunneled traffic already established, the QoS policy is not applied or removed from the traffic stream. To apply or remove the QoS policy for such connections, you must clear (that is, drop) the connections and reestablish them.

Fields

- Enable Priority for this flow—Enables or disables strict scheduling priority for this flow. Priority (LLQ) does not work unless the priority queues are set. To configure priority queues, choose **Configuration > Properties > Priority Queue**. For more information, see Priority Queue.
- Enable policing—Enables policing of traffic in the input or output direction.
  - Direction—Select to enable policing in either the input or output direction.

- 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.

Note The Enable Policing check box merely enforces the maximum speed and burst rate, forcing them to the conforming rate value. It does not enforce the conform-action or the exceed-action specification if these are present.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Edit Class Map

Configuration > Features > Security Policy > Edit > Edit Class Map

The Edit Class Map dialog box lets you add or edit the description of a class map.

Fields

- Description—Add or change the name of the class map description.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Edit Rule

Configuration > Security Policy > Access Rules > Edit Rule

The Edit Rule dialog box lets you modify an existing rule.

Fields

- Select an action—Determines the action type of the new rule. Select either permit or deny from the Select an action drop-down list.
  - Permit—Permits all matching traffic.
  - Deny—Denies all matching traffic.
- Apply to traffic—Determines which type of traffic to which the rule is applied.
  - Incoming to source interface—Selects incoming traffic to the source interface.
  - Outgoing from destination interface—Selects outgoing traffic from the destination interface.
- Syslog—Shows whether syslog is enabled or not.
- More Options—Enables logging for the access list and sets logging options. The More Options button lets you set logging options. This button allows you to:
  - Use default logging behavior.
  - Enable logging for the rule.
  - Disable logging for the rule.
  - Set the level and interval for permit and deny logging. This option checks the Enable Logging check box.

    See Log Options for more information. Also, see Advanced Access Rule Configuration to set global logging options.
- Time Range—Select a time range defined for this rule from the drop-down list.
- New—Create a new time range for this rule. See Add Time Range.
- Source and Destination Host/Network IP Address—Choose this button to identify the networks by IP address.
  - Interface—The interface on which the host or network resides.
  - IP address—The IP address of the host or network.
  - Browse—Lets you select an existing host or network by choosing the options under the **Select Host/Network** window to populate the Name, Interface, IP address, and Mask fields with the properties of the selected host or network.
  - Mask—The subnet mask of the host or network
- Name—Choose this button to identify the networks by name. To name hosts/networks, see the Hosts/Networks tab.

  The name of the host or network. If you choose this option, and reopen the rule to edit it, the button selection reverts to IP Address, and the named host/network IP address information appears in the fields.
- Group—Choose this button to identify a group of networks and hosts that you grouped together on the Hosts/Networks tab.
  - Interface—The interface connected to the hosts and networks in the group.
  - Group—The group name.
- Protocol and Service: TCP and UDP buttons—Selects the TCP/UDP protocol for the rule. The Source Port and Destination Port areas allow you to specify the ports that the access list uses to match packets.
  - Source Port Service—Choose this option to specify a port number, a range of ports, or a well-known service name from a drop-down list of services, such as HTTP or FTP.

  - Source Port Service—The operator drop-down list specifies how the access list matches the port. Choose one of the following operators:
  - = —Equals the port number.
  - not = —Does not equal the port number.
  - > —Greater than the port number.
  - < —Less than the port number.
  - range—Equal to one of the port numbers in the range.
  - Source Port Service—Specifies port number, a range of ports, or a well-known service name from a drop-down list of services, such as HTTP or FTP. The browse button displays the Service dialog box, which lets you select a TCP or UDP service from a preconfigured drop-down list.
  - Source Port Service Group—Choose this option to specify a service group from the Service Group drop-down list,

- Protocol and Service: ICMP—Specifies the ICMP type for the rule in the ICMP type field. The Browse button displays the Service dialog box, which lets you select an ICMP type from a preconfigured drop-down list.
- Protocol and Service: IP—Specifies the IP protocol for the rule in the IP protocol field. The Browse button displays the Protocols dialog box, which lets you select an IP protocol from a preconfigured drop-down list.
- Manage Service Groups—Manages service groups. Service groups allow you to identify multiple non-contiguous port numbers that you want the access list to match. For example, if you want to filter HTTP, FTP, and port numbers 5, 8, and 9, you can define a service group that includes all these ports. Without service groups, you would have to create a separate rule for each port.

  You can create service groups for TCP, UDP, and TCP-UDP. A service group with the TCP-UDP protocol contains services, ports, and ranges that might use either the TCP or UDP protocol. See Manage Service Groups for more information.
- Description—(Optional) Enter a description of the access rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Edit Service Policy Rule > Traffic Classification Tab

Configuration > Security Policy > Edit Service Policy Rule> Traffic Classification Tab

The Traffic Classification tab lets you specify the criteria you want to use to match traffic to which the security policy rule applies.

Fields

- Description—Specifies a description for the traffic classification.

- Default Inspection Traffic—Uses the criteria specified in the default inspection traffic policy.
- Source and destination IP address (uses ACL)—Matches traffic based on the source and destination IP address, using an access control list. This selection is only available if you apply the rule to a specific interface using an Interface Service Policy.
- Tunnel Group—Matches traffic based on the tunnel group.
- TCP or UDP destination port—Matches traffic based on the TCP or UDP destination port.
- RTP range—Matches traffic based on a range of RTP ports.
- IP DiffServ CodePoints (DSCP)—Matches traffic based on the Differentiated Services model of QoS.
- IP Precedence—Matches traffic based on the IP precedence model of QoS.
- Any traffic—Matches all traffic regardless of the traffic type.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Tunnel Group

Add Service Policy Rule Wizard >Traffic Match > Tunnel Group

The Tunnel Group dialog box lets you identify the traffic to which a service policy rule applies based on the tunnel group.

Fields

- Tunnel Group—Selects the tunnel group for which to match traffic.
- New—Displays the Add Tunnel Group window, on which you can configure a new tunnel group.
- Match flow destination IP address—Adds the requirement to match the flow destination IP address, as well as the tunnel group.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

SUNRPC Server

Configuration > Properties > SUNRPC Server

The SUNRPC Server window shows which SunRPC services can traverse the security appliance 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 security appliance.
- 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit SUNRPC Service

Configuration > Properties > SUNRPC Server > Add/Edit SUNRPC Service

The Add/Edit SUNRPC Service dialog box lets you specify what SunRPC services are allowed to traverse the security appliance 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 security appliance.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

CHAPTER 22

NAT

The security appliance supports both the Network Address Translation feature, which provides a globally unique address for each outbound host session, and the Port Address Translation feature, which provides a single, unique global address for up to 64,000 simultaneous outbound or inbound host sessions. The global addresses used for NAT come from a pool of addresses to be used specifically for address translation. The unique global address that is used for PAT can either be one global address or the IP address of a given interface.

The security appliance can perform NAT or PAT in both inbound and outbound connections. This ability to translate inbound addresses is called Outside NAT because addresses on the outside, or less secure, interface are translated to a usable inside IP address. Outside NAT gives you the option to translate an outside host or network to an inside host or network, and it is sometimes referred to as bi-directional NAT. Just as when you translate outbound traffic with NAT, you may choose dynamic NAT, static NAT, dynamic PAT, and static PAT. If necessary, you may use outside NAT together with inside NAT to translate the both source and destination IP addresses of a packet.

NAT

Configuration > Security Policy > NAT

The NAT pane lets you view all the address translation rules or Network Address Translation exemption rules applied to your network.

From the NAT pane you can also create a Translation Exemption Rule, which lets you specify traffic that is exempt from being translated or encrypted. The Exemption Rules are grouped by interface in the table, and then by direction. If you have a group of IP addresses that will be translated, you can exempt certain addresses from being translated using the Exemption Rules. You can use a previously configured access list to define your exemption rule. ASDM will write to the command-line interface a nat 0 command. You can resort the view of your exemption by clicking the column heading.

You can also identify local traffic for address translation by specifying the source and destination addresses (or ports) in an access list using policy NAT. With policy NAT, you can create multiple NAT or static statements that identify the same local address as long as the source/port and destination/port combination is unique for each statement. You can then match different global addresses to each source/port and destination/port pair.

Prerequisites

- Before you can designate access and translation rules for your network, you must first define each host or server for which a rule will apply.

Caution Review Important Notes about Object Groups regarding the naming of Network and Service Groups.

Restrictions

- You cannot use unavailable translation commands until you define networks or hosts. Unavailable commands appear dimmed on the menu.
- It is important to note that the order in which you apply translation rules can affect the way the rules operate. ASDM will list the static translations first and then the dynamic translations. When processing NAT, the security appliance will first translate the static translations in the order they are configured. You can use Insert or Insert After to determine the order in which static translations are processed. Because dynamically translated rules are processed on a best match basis, the option to insert a rule before or after a dynamic translation is disabled.
- It is necessary to run NAT even if you have routable IP addresses on your secure networks. When running NAT with routable IP addresses, translate the routable IP address to itself on the outside.
- A packet sourced on the more secure (inside) interface destined for an intermediate (DMZ) interface can not have the same translated address when it is outbound on a less secure (outside) interface. Furthermore, if one dynamic rule is deleted on either of the outbound interfaces, all outbound dynamic rules for translations originating on the same interface will be deleted.
- It is possible to create an Exemption Rule for traffic so that traffic is sent out to the Internet or a less secure interface unencrypted. This can be useful in a scenario where you want to encrypt some traffic to another remote VPN network, but would like traffic destined to anywhere else to remain unencrypted.

Fields

- Add—Adds a new NAT rule. Choose the type of rule you want to add from the drop-down list.
- Edit—Edits an NAT rule.
- Delete—Deletes a NAT rule.
- Move Up—Moves a rule up. Rules are assessed in the order they appear in this table, so the order can matter if you have overlapping rules.
- Move Down—Moves a rule down.
- Cut—Cuts a rule.
- Copy—Copies the parameters of a rule so you can start a new rule with the same parameters using the Paste button.
- Paste—Opens an Add/Edit Rule dialog box with the copied or cut parameters of the rule prefilled. You can then make any modifications and add it to the table. The Paste button adds the rule above the selected rule. The Paste After item, available from the Paste drop-down list, adds the rule after the selected rule.
- Find—Filters the display to show only matching rules. Clicking **Find** opens the Filter field. Click **Find** again to hide the Filter field.
  - Filter drop-down list—Choose the criteria to filter on, either Interface, Source, Destination, Service, Action, or Rule Query. A rule query is a collection of multiple criteria that you can save and use repeatedly.
  - Filter field—For the Interface type, this field becomes a drop-down list so you can choose an interface name. For the Action type, the drop-down list includes Exempt, Static, and Dynamic. For the Rule Query type, the drop-down list includes all defined rule queries. The Source and Destination types accept an IP address. You can type one manually, or browse for one by

clicking the **...** button and launching the Browse Address dialog box. The Service type accepts a TCP, UDP, TCP-UDP, ICMP, or IP protocol type. You can type one manually, or browse for one by clicking the **...** button and launching the Browse Service Groups dialog box.

  - Filter—Runs the filter.
  - Clear—Clears the Filter field.
  - Rule Query—Opens the Rules Queries dialog box so you can manage named rule queries.
- Show Rule Flow Diagram—Shows the Rule Flow Diagram area under the rule table. This diagram shows the networks, type of traffic, interface name, direction of flow, and action.
- Packet Trace—Opens the Packet Tracer tool with the parameters pre-filled with the characteristics of the selected rule.

The following description summarizes the columns in the NAT Rules table. You can edit the contents of these columns by double-clicking on a table cell. Double-clicking on a column header sorts the table in ascending alphanumeric order, using the selected column as the sort key. If you right-click a rule, you see all of the options represented by the buttons above, as well as Insert and Insert After items. These items either insert a new rule before the selected rule (Insert) or after the selected rule (Insert After.)

- No—Indicates the order of evaluation for the rule.
- Type—Displays the translation rule type applied to the given row, which can either be dynamic or static.
  - Dynamic—Internal IP addresses are dynamically translated using IP addresses from a pool of global addresses or, in the case of PAT, a single address. These rules translate addresses of hosts on a higher security level interface to addresses selected from a pool of addresses for traffic sent to a lower security level interface. Dynamic translations are often used to map local, RFC 1918 IP addresses to addresses that are Internet-routable addresses. They are represented with the dynamic icon.
  - Static—Internal IP addresses are permanently mapped to a global IP address. These rules map a host address on a lower security level interface to a global address on a higher security level interface. For example, a static rule would be used for mapping the local address of a web server on a perimeter network to a global address that hosts on the outside interface would use to access the web server. They are represented with the static icon.
- Real—Displays the original address with its associated interface before network translation is applied.
  - Source Network—The source network on which the traffic to be translated resides for policy NAT. For regular NAT, this displays any.
  - Destination Network—The destination network on which the traffic to be translated resides for policy NAT. For regular NAT, this displays any.
- Translated—Displays the translated addresses and the associated interfaces after network translation is applied.
  - Interface—The interface on which the translated addresses reside.
  - Address—The translated addresses.
- Options—Includes the following items:
  - DNS Rewrite—Lets the security appliance rewrite the DNS record so that an outside client can resolve name of an inside host using an inside DNS server, or vice versa. For example, assume an inside web server www.example.com has IP 192.168.1.1, it is translated to 10.1.1.1 on the outside interface. An outside client sends a DNS request to an inside DNS server, which will

resolve www.example.com to 192.168.1.1. When the reply comes to the security appliance with DNS Rewrite enabled, the security appliance will translate the IP address in the payload to 10.1.1.1, so that the outside client will get the correct IP address.

  - Maximum TCP Connections—The maximum number of TCP connections that are allowed to connect to the statically translated IP address. Valid options are 0 through 65,535. If this value is set to zero, the number of connections is unlimited.
  - Embryonic Limit—The number of embryonic connections allowed to form before the security appliance begins to deny these connections. Set this limit to prevent attack by a flood of embryonic connections. An embryonic connection is one that has been started but has not yet been established, such as a three-way TCP handshake state. Valid values are 0 through 65,535. If this value is set to zero, the number of connections is unlimited. A positive number enables the TCP Intercept feature.
  - Maximum UDP Connections—The maximum number of UDP connections that are allowed to connect to the statically translated IP address. Valid options are 0 through 65,535. If this value is set to zero, the number of connections is unlimited.
  - Randomize Sequence Number—With this check box checked, the security appliance will randomize the sequence number of TCP packets. Disable this feature only if another inline security appliance is also randomizing sequence numbers and the result is scrambling the data. Use of this option opens a security hole in the security appliance. The default is selected.

- Description (for Policy NAT only)—If a description of the rule is available, it is displayed in this column.
- Enable traffic through the firewall without address translation—Allows traffic to pass through the security appliance without address translation.
- Addresses—Tab that lets you add, edit, delete, or find IP address objects, IP names, or network object groups.
- Services—Tab that lets you add, edit, delete, or find services.
- Global Pools—Tab that lets you manage the Global address NAT pools, which are used for dynamic NAT configuration. These are the IP addresses the security appliance will present to the outside or less secure interface for which they are configured.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

Add/Edit Static NAT Rule

Configuration > Security Policy > NAT > Add/Edit Static NAT Rule

The Add/Edit Static NAT Rule dialog box lets you add, edit, and paste translation rules for your security appliance, which are viewed in the NAT Rules table. A Static NAT rule specifies that the address translation is a static, one-to-one translation of an IP address from a private (non-valid) IP address to a global (valid) IP address. Static or Dynamic can be selected, but not both.

Note Review Important Notes about Object Groups regarding the naming of Network and Service Groups.

Fields

- Real Address—The original address with its associated interface before network translation is applied.
  - Interface—Selects the security appliance network interface on which the original host or network resides.
  - IP address—Specifies the IP address of the host or network to which you would like to apply a rule.
  - Mask—Select the network mask (netmask) for the address.
  - Browse—Lets you select the correct IP address and mask from the Hosts/Networks tree for a predefined host or network.
- Static Translation—Lets you specify the static interface and IP address.
  - Interface—Selects the security appliance network interface for static translation.
  - IP address—Selects the IP address for the static translation.
  - Browse—Lets you select the correct IP address and mask from the Hosts/Networks tree from a predefined host or network.
- Enable Port Address Translation (PAT)—Choose this option to specify the protocol, original port, and translated port for PAT.
  - Protocol—TCP or UDP.
  - Original Port—Select from the list of ports.
  - Translated Port—Select from the list of ports.
- NAT Options—Lets you configure the DNS Rewrite, Maximum Connections, Embryonic Limit and Randomize Sequence Number.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

Add/Edit Dynamic NAT Rule

Configuration > Security Policy > NAT > Add/Edit Dynamic NAT Rule

The Add/Edit Dynamic NAT Rule dialog box lets you add, edit, and paste translation rules for your security appliance, which are viewed in the NAT Rules table. A Dynamic NAT rule specifies either a predefined pool of IP addresses, or to perform PAT on a global IP address or the less secure interface for multiple hosts on the more secure interface. For example, if your inside network has multiple hosts, you

can permit outbound access through a pool or a PAT address by using Dynamic NAT to dynamically assign an global IP address for each host requesting an outbound connection. Static or Dynamic can be selected, but not both.

Note Review Important Notes about Object Groups regarding the naming of Network and Service Groups.

Fields

- Real Address—The original address with its associated interface before network translation is applied.
  - Interface—Selects the security appliance network interface on which the original host or network resides.
  - IP Address—Specifies the IP address of the host or network to which you would like to apply a rule.
  - Mask—Select the network mask (netmask) for the address.
  - Browse—Lets you select the correct IP address and mask from the Hosts/Networks tree for a predefined host or network.
- Dynamic Translation—Lets you specify the dynamic interface and global address pool.
  - Interface—Selects the security appliance network interface for dynamic translation.
  - Add—Adds a global pool.
  - Edit—Edits a global pool.
  - Delete—Deletes a global pool.
- NAT Options—Lets you configure the DNS Rewrite, Maximum Connections, Embryonic Limit and Randomize Sequence Number.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

NAT Options

Configuration > Security Policy > NAT > Add/Edit NAT Rule> NAT Options

The NAT Options dialog box lets you configure the DNS Rewrite, Maximum Connections, Embryonic Limit, and Randomize Sequence Number for NAT and Policy NAT.

Fields

- DNS Rewrite—Lets the security appliance rewrite the DNS record so that an outside client can resolve name of an inside host using an inside DNS server, or vice versa. For example, assume an inside web server www.example.com has IP 192.168.1.1, it is translated to 10.1.1.1 on the outside interface. An outside client sends a DNS request to an inside DNS server, which will resolve

www.example.com to 192.168.1.1. When the reply comes to the security appliance with DNS Rewrite enabled, the security appliance will translate the IP address in the payload to 10.1.1.1, so that the outside client will get the correct IP address.

- Maximum TCP Connections—The maximum number of TCP connections that are allowed to connect to the statically translated IP address. Valid options are 0 through 65,535. If this value is set to zero, the number of connections is unlimited.
- Maximum UDP Connections—The maximum number of UDP connections that are allowed to connect to the statically translated IP address. Valid options are 0 through 65,535. If this value is set to zero, the number of connections is unlimited.
- Embryonic Limit—The number of embryonic connections allowed to form before the security appliance begins to deny these connections. Set this limit to prevent attack by a flood of embryonic connections. An embryonic connection is one that has been started but has not yet been established, such as a three-way TCP handshake state. Valid values are 0 through 65,535. If this value is set to zero, the number of connections is unlimited. A positive number enables the TCP Intercept feature.
- Randomize Sequence Number—With this check box checked, the security appliance will randomize the sequence number of TCP packets. Disable this feature only if another inline security appliance is also randomizing sequence numbers and the result is scrambling the data. Use of this option opens a security hole in the security appliance. The default is selected.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

Global Pools

Configuration > Security Policy > NAT > Add/Edit Address Translation Rule > Global Pools

The Global Pools dialog box lets you view, define new, or delete existing global address pools used in dynamic NAT rules. For more information on dynamic NAT rules and its uses, refer to Understanding Dynamic NAT.

Fields

- Interface—Identifies the interface name associated with the address pool used for dynamic address translation.
- Pool ID—Identifies the ID number of the address pool.
- IP Address(es)—Identifies the type and value of the address(es) for the pool. It can identify one of the following types:
  - A range of addresses
  - A PAT address
  - A PAT address associated with an interface

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

Add/Edit Global Address Pool

Configuration > Security Policy > NAT > Add/Edit Address Translation Rule > Global Pools > Add/Edit Global Address Pool

The **Add/Edit Global Address Pool** dialog box lets you define the settings for a new global address pool or edit the settings of an existing pool.

Fields

- Interface—Specifies the interface name to associate with the new address pool. Select the name in the Interface drop-down list.
- Pool ID—Specifies the ID number that dynamic NAT rules use to reference this address pool. Enter the number in the Pool ID field.
- Range—Select this option to specify that a range of IP addresses be used with the new address pool. If you select this option, specify the following values:
  - Enter the start and end addresses used by the range in the **IP Address** fields. These addresses are the addresses to which the original addresses will be translated. If the security appliance is exposing the host or network to users on the Internet, these IP addresses must be valid IP addresses that are registered with the American Registry for Internet Numbers.
  - Enter the mask in the **Network Mask** (optional) field. This value identifies the mask of the network on which translated IP addresses are members.
- Port Address Translation (PAT)—Choose this option to specify that an IP address be used for Port Address Translation. If you select this option, specify the following value:
  - Enter the IP address used for PAT in the **IP Address** field.
    This value is the specific translated IP address to which you want to translate the original addresses of the translated host or network. If the security appliance is exposing the host or network to users on the Internet, this IP address must be a valid IP address that is registered with ARIN.
- Port Address Translation (PAT) using the IP address of the interface—Select this option to specify that the IP address assigned to the interface selected in the **Interface** drop-down list be used as the translated address for PAT.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

Add/Edit Static Policy NAT Rule

Configuration > Security Policy > NAT > Add/Edit Static Policy NAT Rule

The Add/Edit Static Policy NAT Rule dialog box lets you configure the protocol and service that policy NAT will use to translate traffic.

Fields

- Real Address—The original address with its associated interface before network translation is applied.
  - Interface—Selects the security appliance network interface on which the original host or network resides.
  - Source—Choose type, IP address, and netmask.
  - Destination—Choose type, IP address, and netmask
- Protocol and Service—Lets you define the protocols and services to be used for policy NAT.
  - TCP—Select to define the TCP protocol types used for translation with policy NAT.
  - UDP—Select to define the UDP protocol types used for translation with policy NAT.
  - ICMP—Select to define the ICMP protocol types used for translation with policy NAT.
  - IP—Select to define the IP protocol types used for translation with policy NAT.
  - IP Protocol—Depending on your selection this displays the TCP, UDP, ICMP or IP protocol type. You can either enter the port or protocol number or select the protocol from a drop-down list using the browse (...) button.
- Static Translation—Lets you specify the static interface and IP address.
  - Interface—Selects the security appliance network interface for static translation.
  - IP address—Selects the IP address for the static translation.
  - Browse—Lets you select the correct IP address and mask from the Hosts/Networks tree from a predefined host or network.
- Enable Port Address Translation (PAT)—Choose this option to specify the protocol, original port, and translated port for PAT.
  - Protocol—TCP or UDP.
  - Original Port—Select from the list of ports.
  - Translated Port—Select from the list of ports.
- NAT Options—Lets you configure the DNS Rewrite, Maximum Connections, Embryonic Limit and Randomize Sequence Number.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

Add/Edit Dynamic Policy NAT Rule

Configuration > Security Policy > NAT > Add/Edit Dynamic Policy NAT Rule

The Add/Edit Dynamic Policy NAT Rule dialog box lets you configure the protocol and service that policy NAT will use to translate traffic.

Fields

- Real Address—The original address with its associated interface before network translation is applied.
  - Interface—Selects the security appliance network interface on which the original host or network resides.
  - Source—Choose type, IP address, and netmask.
  - Destination—Choose type, IP address, and netmask
- Protocol and Service—Lets you define the protocols and services to be used for policy NAT.
  - TCP—Select to define the TCP protocol types used for translation with policy NAT.
  - UDP—Select to define the UDP protocol types used for translation with policy NAT.
  - ICMP—Select to define the ICMP protocol types used for translation with policy NAT.
  - IP—Select to define the IP protocol types used for translation with policy NAT.
  - IP Protocol—Depending on your selection this displays the TCP, UDP, ICMP or IP protocol type. You can either enter the port or protocol number or select the protocol from a drop-down list using the browse (...) button.
- Dynamic Translation—Lets you specify the dynamic interface and global address pool.
- Real Address—The original address with its associated interface before network translation is applied.
  - Interface—Selects the security appliance network interface on which the original host or network resides.
  - IP Address—Specifies the IP address of the host or network to which you would like to apply a rule.
  - Mask—Select the network mask (netmask) for the address.
  - Browse—Lets you select the correct IP address and mask from the Hosts/Networks tree for a predefined host or network.
- Dynamic Translation—Lets you specify the dynamic interface and global address pool.
  - Interface—Selects the security appliance network interface for dynamic translation.
  - Add—Adds a global pool.
  - Edit—Edits a global pool.

  - Delete—Deletes a global pool.
- NAT Options—Lets you configure the DNS Rewrite, Maximum Connections, Embryonic Limit and Randomize Sequence Number.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

Add/Edit NAT Exempt Rule

Configuration > Security Policy > NAT > Add/Edit NAT Exempt Rule

The Add/Edit NAT Exempt Rule dialog box lets you add and edit Network Address Translation exemption rules for your security appliance. Depending upon which command you selected in the Translation Rules menu, the title for this dialog box will appear as Add Address Exemption Rule or Edit Address Exemption Rule.

Fields

- Action—The action drop-down list lets you select the action, exempt or do not exempt, that this exemption rule will take if the host/network meets the criteria defined. The Select an action list options are as follows:
  - Exempt—Specifies that the traffic defined will be exempted from NAT.
  - Do Not Exempt—Specifies that the traffic defined will not be exempted from NAT.
- IP Address—Selects the criteria of testing the IP address of the source host or network to determine if the action of the exemption rule will be applied. Selecting this option displays the following fields:
  - Interface—Selects the security appliance network interface name on which the original host or network resides.
  - IP address—Specifies the IP address of the host or network to which you would like to apply a rule.
  - Browse—Lets you select the correct IP address and mask from the Hosts/Networks tree from a predefined host or network.
  - Mask—Select the network mask (netmask) for the address.
- Name—Selects the criteria of testing the name of the source host or network to determine if the action of the exemption rule will be applied. Selecting this option displays the following fields:
  - Name—Lets you select a previously defined name of a host or network to which you would like to apply the rule. The security appliance also automatically generates a hostname for each interface by using the interface name, such as inside or outside.
- Group—Selects the criteria of testing a group of the source host or network to determine if the action of the exemption rule will be applied. Selecting this option displays the following fields:
  - Interface—Selects the security appliance network interface name on which the original host or network resides.

  - Group—Selects the group of the host or network to which you would like to apply the rule.
- When Connecting To—The When Connecting To area lets you define the criteria which must be met for the action to be performed. The criteria my be defined by selecting an IP address, Name, Group, or by browsing a previously defined drop-down list of hosts/networks.
- IP address—Specifies the IP address of the destination host or network to which you would like to apply the exemption rule.
  - Interface—Selects the security appliance network interface name on which the original host or network resides.
  - IP address—Specifies the IP address of the host or network to which you would like to apply a rule.
  - Browse—Lets you select the correct IP address and mask from the Hosts/Networks tree from a predefined host or network.
  - Mask—Select the network mask (netmask) for the address.
- Name—Selects the criteria of testing the name of the source host or network to determine if the action of the exemption rule will be applied. Selecting this option displays the following fields:
  - Name—Lets you select a previously defined name of a host or network to which you would like to apply the rule. The security appliance also automatically generates a hostname for each interface by using the interface name, such as inside or outside.
- Group—Selects the criteria of testing a group of the source host or network to determine if the action of the exemption rule will be applied. Selecting this option displays the following fields:
  - Interface—Selects the security appliance network interface name on which the original host or network resides.
  - Group—Selects the group of the host or network to which you would like to apply the rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

Add/Edit Identity NAT Rule

Configuration > Security Policy > NAT > Add/Edit Identity NAT Rule

The Add/Edit Identity NAT Rule dialog box lets you configure the identity NAT settings.

Fields

- Real Address—The original address with its associated interface before network translation is applied.
  - Interface—Selects the security appliance network interface on which the original host or network resides.

  - IP address—Specifies the IP address of the host or network to which you would like to apply a rule.
  - Mask—Select the network mask (netmask) for the address.
  - Browse—Lets you select the correct IP address and mask from the Hosts/Networks tree for a predefined host or network.
- Enable outside NAT—Choose this option to enable outside NAT.
- NAT Options—Lets you configure the DNS Rewrite, Maximum Connections, Embryonic Limit and Randomize Sequence Number.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

Configuring ARP Inspection and Bridging Parameters

This chapter describes how to enable ARP inspection and how to customize bridging operations for the security appliance in transparent firewall mode. In multiple context mode, the commands in this chapter can be entered in a security context, but not the system.

For information about transparent firewall mode, see Chapter 16, "Firewall Mode Overview."

This chapter includes the following sections:

Configuring ARP Inspection

This section describes ARP inspection and how to enable it, and includes the following topics:

ARP Inspection

Configuration > Properties > ARP > ARP Inspection

The ARP Inspection pane lets you configure ARP inspection.

By default, all ARP packets are allowed through the security appliance. You can control the flow of ARP packets by enabling ARP inspection.

When you enable ARP inspection, the security appliance compares the MAC address, IP address, and source interface in all ARP packets to static entries in the ARP table, and takes the following actions:

- If the IP address, MAC address, and source interface match an ARP entry, the packet is passed through.
- If there is a mismatch between the MAC address, the IP address, or the interface, then the security appliance drops the packet.

- If the ARP packet does not match any entries in the static ARP table, then you can set the security appliance to either forward the packet out all interfaces (flood), or to drop the packet.

Note The dedicated management interface, if present, never floods packets even if this parameter is set to flood.

ARP inspection prevents malicious users from impersonating other hosts or routers (known as ARP spoofing). ARP spoofing can enable a "man-in-the-middle" attack. For example, a host sends an ARP request to the gateway router; the gateway router responds with the gateway router MAC address. The attacker, however, sends another ARP response to the host with the attacker MAC address instead of the router MAC address. The attacker can now intercept all the host traffic before forwarding it on to the router.

ARP inspection ensures that an attacker cannot send an ARP response with the attacker MAC address, so long as the correct MAC address and the associated IP address are in the static ARP table.

Fields

- Interface—Shows the interface names.
- ARP Inspection Enabled—Shows if ARP inspection is enabled, Yes or No.
- Flood Enabled—If ARP inspection is enabled, shows if the action is to flood unknown packets, Yes or No. If ARP inspection is disabled, this value is always No.
- Edit—Edits the ARP inspection parameters for the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
—	•	•	•	—

Edit ARP Inspection Entry

Configuration > Properties > ARP > ARP Inspection > Edit ARP Inspection Entry

The Edit ARP Inspection Entry dialog box lets you set ARP inspection settings.

Fields

- Enable ARP Inspection—Enables ARP inspection.
- Flood ARP Packets—Specifies that packets that do not match any element of a static ARP entry are flooded out all interfaces except the originating interface. If there is a mismatch between the MAC address, the IP address, or the interface, then the security appliance drops the packet. If you do not check this check box, all non-matching packets are dropped.

Note The default setting is to flood non-matching packets. To restrict ARP through the security appliance to only static entries, then set this command to **no-flood**.

The Management 0/0 interface or subinterface, if present, never floods packets even if this parameter is set to flood.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
—	•	•	•	—

ARP Static Table

Configuration > Properties > ARP > ARP Static Table

Although hosts identify a packet destination by an IP address, the actual delivery of the packet on Ethernet relies on the Ethernet MAC address. When a router or host wants to deliver a packet on a directly connected network, it sends an ARP request asking for the MAC address associated with the IP address, and then delivers the packet to the MAC address according to the ARP response. The host or router keeps an ARP table so it does not have to send ARP requests for every packet it needs to deliver. The ARP table is dynamically updated whenever ARP responses are sent on the network, and if an entry is not used for a period of time, it times out. If an entry is incorrect (for example, the MAC address changes for a given IP address), the entry times out before it can be updated.

Note The transparent firewall uses dynamic ARP entries in the ARP table for traffic to and from the security appliance, such as management traffic.

The ARP Static Table panel lets you add static ARP entries that map a MAC address to an IP address for a given interface. Static ARP entries do not time out, and might help you solve a networking problem.

Fields

- Interface—Shows the interface attached to the host network.
- IP Address—Shows the host IP address.
- MAC Address—Shows the host MAC address.
- Proxy ARP—Shows whether the security appliance performs proxy ARP for this address. If the security appliance receives an ARP request for the specified IP address, then it responds with the specified MAC address.
- Add—Adds a static ARP entry.
- Edit—Edits a static ARP entry.
- Delete—Deletes a static ARP entry.

- ARP Timeout—Sets the amount of time before the security appliance rebuilds the ARP table, between 60 to 4294967 seconds. The default is 14400 seconds. Rebuilding the ARP table automatically updates new host information and removes old host information. You might want to reduce the timeout because the host information changes frequently. Although this parameter appears on the ARP Static Table panel, the timeout applies to the *dynamic* ARP table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit ARP Static Configuration

Configuration > Properties > ARP > ARP Static Table > Add/Edit ARP Static Configuration

The Add/Edit ARP Static Configuration dialog box lets you add or edit a static ARP entry.

Fields

- Interface—Sets the interface attached to the host network.
- IP Address—Sets the host IP address.
- MAC Address—Sets the host MAC address; for example, 00e0.1e4e.3d8b.
- Proxy ARP—Enables the security appliance to perform proxy ARP for this address. If the security appliance receives an ARP request for the specified IP address, then it responds with the specified MAC address.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Customizing the MAC Address Table

This section describes the MAC address table, and includes the following topics:

MAC Address Table

Configuration > Properties > Bridging > MAC Address Table

The MAC Address Table pane lets you add static MAC Address entries. Normally, MAC addresses are added to the MAC address table dynamically as traffic from a particular MAC address enters an interface. You can add static MAC addresses to the MAC address table if desired. One benefit to adding static entries is to guard against MAC spoofing. If a client with the same MAC address as a static entry attempts to send traffic to an interface that does not match the static entry, then the security appliance drops the traffic and generates a system message. When you add a static ARP entry (see the "ARP Static Table" section on page 23-3), a static MAC address entry is automatically added to the MAC address table.

The security appliance learns and builds a MAC address table in a similar way as a normal bridge or switch: when a device sends a packet through the security appliance, the security appliance adds the MAC address to its table. The table associates the MAC address with the source interface so that the security appliance knows to send any packets addressed to the device out the correct interface.

The ASA 5505 adaptive security appliance includes a built-in switch; the switch MAC address table maintains the MAC address-to-switch port mapping for traffic within each VLAN. This section discusses the bridge MAC address table, which maintains the MAC address-to-VLAN interface mapping for traffic that passes between VLANs.

Because the security appliance is a firewall, if the destination MAC address of a packet is not in the table, the security appliance does not flood the original packet on all interfaces as a normal bridge does. Instead, it generates the following packets for directly connected devices or for remote devices:

- Packets for directly connected devices—The security appliance generates an ARP request for the destination IP address, so that the security appliance can learn which interface receives the ARP response.
- Packets for remote devices—The security appliance generates a ping to the destination IP address so that the security appliance can learn which interface receives the ping reply.

The original packet is dropped.

Fields

- Interface—Shows the interface associated with the MAC address.
- MAC Address—Shows the MAC address.
- Add—Adds a static MAC address entry.
- Edit—Edits a static MAC address entry.
- Delete—Deletes a static MAC address entry.
- Dynamic Entry Timeout—Sets the time a MAC address entry stays in the MAC address table before timing out, between 5 and 720 minutes (12 hours). 5 minutes is the default.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
—	•	•	•	—

Add/Edit MAC Address Entry

Configuration > Properties > Bridging > MAC Address Table > Add/Edit MAC Address Entry

The Add/Edit MAC Address Entry dialog box lets you add or edit a static MAC address entry. Normally, MAC addresses are added to the MAC address table dynamically as traffic from a particular MAC address enters an interface. One benefit to adding static entries is to guard against MAC spoofing. If a client with the same MAC address as a static entry attempts to send traffic to an interface that does not match the static entry, then the security appliance drops the traffic and generates a system message.

Fields

- Interface Name—Sets the interface associated with the MAC address.
- MAC Address—Sets the MAC address.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
—	•	•	•	—

MAC Learning

Configuration > Properties > Bridging > MAC Learning

The MAC Learning pane lets you disable MAC address learning on an interface. By default, each interface automatically learns the MAC addresses of entering traffic, and the security appliance adds corresponding entries to the MAC address table. You can disable MAC address learning if desired; however, unless you statically add MAC addresses to the table, no traffic can pass through the security appliance.

Fields

- Interface—Shows the interface name.
- MAC Learning Enabled—Shows if MAC learning is enabled, Yes or No.
- Enable—Enables MAC learning to the selected interface.
- Disable—Disables MAC learning to the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
—	•	•	•	—

Preventing Network Attacks

This chapter describes how to prevent network attacks by configuring protection features, and includes the following sections:

Connection Settings (Transparent Mode Only)

Configuration > Properties > Connection Settings

The Connection Settings pane lets you set maximum TCP and UDP connections, maximum embryonic connections, and lets you disable TCP sequence randomization for outbound traffic (inside to outside) in transparent firewall mode.

Note You can also configure maximum connections, maximum embryonic connections, and TCP sequence randomization in Service Policy Rules. Service Policy Rules provide greater control of the application of these limits, and you can configure limits for traffic in both directions, not just outbound connections. If you configure these settings for the same traffic using both methods, then the security appliance uses the lower limit. For TCP sequence randomization, if it is disabled using either method, then the security appliance disables TCP sequence randomization.

Limiting the number of connections and embryonic connections protects you from a DoS attack. The security appliance uses the embryonic 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 sequence randomization should only be disabled if another in-line firewall is also randomizing sequence numbers and the result is scrambling the data. Each TCP connection has two Initial Sequence Numbers (ISNs): one generated by the client and one generated by the server. The security appliance randomizes the ISN that is generated by the host/server. At least one of the ISNs must be randomly generated so that attackers cannot predict the next ISN and potentially hijack the session.

Fields

- Interface—Shows the Interface on which connection limits are enabled. This interface is always the inside interface, because connection limits are not supported on the outside interface.
- Address—Shows the addresses for which you apply connection limits.
- Maximum TCP Connections—Shows the maximum TCP connections. A value of 0 means unlimited connections.
- Embryonic Limit—Shows the maximum embryonic connections. A value of 0 means unlimited connections.
- Maximum UDP Connections—Shows the maximum UDP connections. A value of 0 means unlimited connections.
- Randomize Sequence Number—Shows if TCP sequence randomization is enabled or disabled, Yes or No.
- Add—Adds a connection limit rule.
- Edit—Edits a connection limit rule.
- Delete—Deletes a connection limit rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
—	•	•	•	—

Set/Edit Connection Settings

Configuration > Properties > Connection Settings > Set/Edit Connection Settings

The Set/Edit Connection Settings dialog box lets you define a connection limit rule for outbound traffic (inside to outside) in transparent firewall mode.

Fields

- Host/Network—Sets the hosts and networks for which you want to set connection limits.
  - Interface—Sets the interface on which you want to set connection limits. Choose the inside interface only.
  - IP Address—Sets the IP addresses for which you want to set connection limits.
  - Mask—Sets the subnet mask. You can either enter a mask in the field, or choose a common mask from the list.
  - Browse—Opens the Select host/network dialog box from which you can choose hosts and networks you defined in the Network Object Groups panel.
- Maximum Connections—Sets the maximum TCP and UDP connections.
  - Maximum TCP Connections—Sets the maximum TCP connections, between 0 and 65535. The 0 value means unlimited connections.

  - Maximum UDP Connections—Sets the maximum UDP connections, between 0 and 65535. The 0 value means unlimited connections.

- Maximum Embryonic Connections—Sets the maximum embryonic connections, between 0 and 65535. The 0 value means unlimited connections. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. The security appliance uses the embryonic limit to trigger TCP Intercept, which protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets.
- Randomize Sequence Number check—Enables TCP sequence number randomization. Uncheck this box to disable randomiztion. TCP sequence randomization should only be disabled if another in-line firewall is also randomizing sequence numbers and the result is scrambling the data.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
—	•	•	•	—

IP Audit

The IP audit feature provides basic IPS functionality; for advanced IPS functionality on supported platforms, you can install an AIP SSM.

This feature lets you create a named audit policy that identifies the actions to take when a packet matches a predefined attack signature or informational signature. Signatures are activities that match known attack patterns. For example, there are signatures that match DoS attacks. You can configure the security appliance to drop the packet, generate an alarm, or reset the connection.

IP Audit Policy

Configuration > Properties > IP Audit > IP Audit Policy

The IP Audit Policy pane lets you add audit policies and assign them to interfaces. You can assign an attack policy and an informational policy to each interface. The attack policy determines the action to take with packets that match an attack signature; the packet might be part of an attack on your network, such as a DoS attack. The informational policy determines the action to take with packets that match an informational signature; the packet is not currently attacking your network, but could be part of an information-gathering activity, such as a port sweep. For a complete list of signatures, see the IP Audit Signature List.

Fields

- Name—Shows the names of the defined IP audit policies. Although the default actions for a named policy are listed in this table ("--Default Action--"), they are not named policies that you can assign to an interface. Default actions are used by named policies if you do not set an action for the policy. You can modify the default actions by selecting them and clicking the Edit button.
- Type—Shows the policy type, either Attack or Info.

- Action—Shows the actions taken against packets that match the policy, Alarm, Drop, and/or Reset. Multiple actions can be listed.
- Add—Adds a new IP audit policy.
- Edit—Edits an IP audit policy or the default actions.
- Delete—Deletes an IP audit policy. You cannot delete a default action.
- Policy-to-Interface Mappings—Assigns an attack and informational policy to each interface.
  - Interface—Shows the interface name.
  - Attack Policy—Lists the attack audit policy names available. Assign a policy to an interface by clicking the name in the list.
  - Info Policy—Lists the informational audit policy names available. Assign a policy to an interface by clicking the name in the list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit IP Audit Policy Configuration

Configuration > Properties > IP Audit > IP Audit Policy > Add/Edit IP Audit Policy Configuration

The Add/Edit IP Audit Policy Configuration dialog box lets you add or edit a named IP audit policy that you can assign to interfaces, and lets you modify the default actions for each signature type.

Fields

- Policy Name—Sets the IP audit policy name. You cannot edit the name after you add it.
- Policy Type—Sets the policy type. You cannot edit the policy type after you add it.
  - Attack—Sets the policy type as attack.
  - Information—Sets the policy type as informational.
- Action—Sets one or more actions to take when a packet matches a signature. If you do not choose an action, then the default policy is used.
  - Alarm—Generates a system message showing that a packet matched a signature. For a complete list of signatures, see IP Audit Signature List.
  - Drop—Drops the packet.
  - Reset—Drops the packet and closes the connection.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

IP Audit Signatures

Configuration > Properties > IP Audit > IP Audit Signatures

The IP Audit Signatures pane lets you disable audit signatures. You might want to disable a signature if legitimate traffic continually matches a signature, and you are willing to risk disabling the signature to avoid large numbers of alarms.

For a complete list of signatures, see IP Audit Signature List.

Fields

- Enabled—Lists the enabled signatures.
- Disabled—Lists the disabled signatures.
- Disable—Moves the selected signature to the Disabled pane.
- Enable—Moves the selected signature to the Enabled pane.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

IP Audit Signature List

Table 24-1 lists supported signatures and system message numbers.

Table 24-1 Signature IDs and System Message Numbers

Signature ID	Message Number	Signature Title	Signature Type	Description
1000	400000	IP options-Bad Option List	Informational	Triggers on receipt of an IP datagram where the list of IP options in the IP datagram header is incomplete or malformed. The IP options list contains one or more options that perform various network management or debugging tasks.
1001	400001	IP options-Record Packet Route	Informational	Triggers on receipt of an IP datagram where the IP option list for the datagram includes option 7 (Record Packet Route).
1002	400002	IP options-Timestamp	Informational	Triggers on receipt of an IP datagram where the IP option list for the datagram includes option 4 (Timestamp).
1003	400003	IP options-Security	Informational	Triggers on receipt of an IP datagram where the IP option list for the datagram includes option 2 (Security options).
1004	400004	IP options-Loose Source Route	Informational	Triggers on receipt of an IP datagram where the IP option list for the datagram includes option 3 (Loose Source Route).
1005	400005	IP options-SATNET ID	Informational	Triggers on receipt of an IP datagram where the IP option list for the datagram includes option 8 (SATNET stream identifier).
1006	400006	IP options-Strict Source Route	Informational	Triggers on receipt of an IP datagram in which the IP option list for the datagram includes option 2 (Strict Source Routing).
1100	400007	IP Fragment Attack	Attack	Triggers when any IP datagram is received with an offset value less than 5 but greater than 0 indicated in the offset field.
1102	400008	IP Impossible Packet	Attack	Triggers when an IP packet arrives with source equal to destination address. This signature will catch the so-called Land Attack.

Table 24-1 Signature IDs and System Message Numbers (continued)

Signature ID	Message Number	Signature Title	Signature Type	Description
1103	400009	IP Overlapping Fragments (Teardrop)	Attack	Triggers when two fragments contained within the same IP datagram have offsets that indicate that they share positioning within the datagram. This could mean that fragment A is being completely overwritten by fragment B, or that fragment A is partially being overwritten by fragment B. Some operating systems do not properly handle fragments that overlap in this manner and may throw exceptions or behave in other undesirable ways upon receipt of overlapping fragments, which is how the Teardrop attack works to create a DoS.
2000	400010	ICMP Echo Reply	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 0 (Echo Reply).
2001	400011	ICMP Host Unreachable	Informational	Triggers when an IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 3 (Host Unreachable).
2002	400012	ICMP Source Quench	Informational	Triggers when an IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 4 (Source Quench).
2003	400013	ICMP Redirect	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 5 (Redirect).
2004	400014	ICMP Echo Request	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 8 (Echo Request).
2005	400015	ICMP Time Exceeded for a Datagram	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 11(Time Exceeded for a Datagram).
2006	400016	ICMP Parameter Problem on Datagram	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 12 (Parameter Problem on Datagram).
2007	400017	ICMP Timestamp Request	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 13 (Timestamp Request).

Table 24-1 Signature IDs and System Message Numbers (continued)

Signature ID	Message Number	Signature Title	Signature Type	Description
2008	400018	ICMP Timestamp Reply	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 14 (Timestamp Reply).
2009	400019	ICMP Information Request	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 15 (Information Request).
2010	400020	ICMP Information Reply	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 16 (ICMP Information Reply).
2011	400021	ICMP Address Mask Request	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 17 (Address Mask Request).
2012	400022	ICMP Address Mask Reply	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 18 (Address Mask Reply).
2150	400023	Fragmented ICMP Traffic	Attack	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and either the more fragments flag is set to 1 (ICMP) or there is an offset indicated in the offset field.
2151	400024	Large ICMP Traffic	Attack	Triggers when a IP datagram is received with the protocol field of the IP header set to 1(ICMP) and the IP length > 1024.
2154	400025	Ping of Death Attack	Attack	Triggers when a IP datagram is received with the protocol field of the IP header set to 1(ICMP), the Last Fragment bit is set, and (IP offset * 8) + (IP data length) > 65535 that is to say, the IP offset (which represents the starting position of this fragment in the original packet, and which is in 8 byte units) plus the rest of the packet is greater than the maximum size for an IP packet.
3040	400026	TCP NULL flags	Attack	Triggers when a single TCP packet with none of the SYN, FIN, ACK, or RST flags set has been sent to a specific host.
3041	400027	TCP SYN+FIN flags	Attack	Triggers when a single TCP packet with the SYN and FIN flags are set and is sent to a specific host.

Table 24-1 Signature IDs and System Message Numbers (continued)

Signature ID	Message Number	Signature Title	Signature Type	Description
3042	400028	TCP FIN only flags	Attack	Triggers when a single orphaned TCP FIN packet is sent to a privileged port (having port number less than 1024) on a specific host.
3153	400029	FTP Improper Address Specified	Informational	Triggers if a port command is issued with an address that is not the same as the requesting host.
3154	400030	FTP Improper Port Specified	Informational	Triggers if a port command is issued with a data port specified that is <1024 or >65535.
4050	400031	UDP Bomb attack	Attack	Triggers when the UDP length specified is less than the IP length specified. This malformed packet type is associated with a denial of service attempt.
4051	400032	UDP Snork attack	Attack	Triggers when a UDP packet with a source port of either 135, 7, or 19 and a destination port of 135 is detected.
4052	400033	UDP Chargen DoS attack	Attack	This signature triggers when a UDP packet is detected with a source port of 7 and a destination port of 19.
6050	400034	DNS HINFO Request	Informational	Triggers on an attempt to access HINFO records from a DNS server.
6051	400035	DNS Zone Transfer	Informational	Triggers on normal DNS zone transfers, in which the source port is 53.
6052	400036	DNS Zone Transfer from High Port	Informational	Triggers on an illegitimate DNS zone transfer, in which the source port is not equal to 53.
6053	400037	DNS Request for All Records	Attack	Triggers on a DNS request for all records.
6100	400038	RPC Port Registration	Informational	Triggers when attempts are made to register new RPC services on a target host.
6101	400039	RPC Port Unregistration	Informational	Triggers when attempts are made to unregister existing RPC services on a target host.
6102	400040	RPC Dump	Informational	Triggers when an RPC dump request is issued to a target host.
6103	400041	Proxied RPC Request	Attack	Triggers when a proxied RPC request is sent to the portmapper of a target host.
6150	400042	ypserv (YP server daemon) Portmap Request	Informational	Triggers when a request is made to the portmapper for the YP server daemon (ypserv) port.
6151	400043	ypbind (YP bind daemon) Portmap Request	Informational	Triggers when a request is made to the portmapper for the YP bind daemon (ypbind) port.

Table 24-1 Signature IDs and System Message Numbers (continued)

Signature ID	Message Number	Signature Title	Signature Type	Description
6152	400044	yppasswdd (YP password daemon) Portmap Request	Informational	Triggers when a request is made to the portmapper for the YP password daemon (yppasswdd) port.
6153	400045	ypupdated (YP update daemon) Portmap Request	Attack	Triggers when a request is made to the portmapper for the YP update daemon (ypupdated) port.
6154	400046	ypxfrd (YP transfer daemon) Portmap Request	Attack	Triggers when a request is made to the portmapper for the YP transfer daemon (ypxfrd) port.
6155	400047	mountd (mount daemon) Portmap Request	Informational	Triggers when a request is made to the portmapper for the mount daemon (mountd) port.
6175	400048	rexd (remote execution daemon) Portmap Request	Informational	Triggers when a request is made to the portmapper for the remote execution daemon (rexd) port.
6180	400049	rexd (remote execution daemon) Attempt	Informational	Triggers when a call to the rexd program is made. The remote execution daemon is the server responsible for remote program execution. This may be indicative of an attempt to gain unauthorized access to system resources.
6190	400050	statd Buffer Overflow	Attack	Triggers when a large statd request is sent. This could be an attempt to overflow a buffer and gain access to system resources.

Fragment

Configuration > Properties > Advanced > Fragment

The Fragment pane lets you configure the IP fragment database on each interface of the security appliance to improve compatibility with NFS.

Fields

- Fragment table:
  - Interface—Lists the available interfaces of the security appliance.
  - Size—Sets the maximum number of packets that can be in the IP reassembly database waiting for reassembly. The default is 200.
  - Chain Length—Specifies the maximum number of packets into which a full IP packet can be fragmented. The default is 24 packets.
  - Timeout—Specifies the maximum number of seconds to wait for an entire fragmented packet to arrive. The timer starts after the first fragment of a packet arrives. If all fragments of the packet do not arrive by the number of seconds specified, all fragments of the packet that were already received will be discarded. The default is 5 seconds.

- Edit—Opens the Edit Fragment dialog box.
- Show Fragment—Opens a panel and displays the current IP fragment database statistics for each interface of the security appliance.

Changing Fragment Parameters

To modify the IP fragment database parameters of an interface, perform the following steps:

Step 1 Choose the interface to change in the Fragment table and click **Edit**. The Edit Fragment dialog box appears.

Step 2 In the Edit Fragment dialog box, change the Size, Chain, and Timeout values as desired, and click **OK**. If you make a mistake, click **Restore Defaults**.

Step 3 Click **Apply** in the Fragment panel.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Show Fragment

Configuration > Properties > Fragment > Show Fragment

The Show Fragment panel displays the operational data of the IP fragment reassembly module.

Fields

- Size—*Display only.* Displays the number of packets in the IP reassembly database waiting for reassembly. The default is 200.
- Chain—*Display only.* Displays the number of packets into which a full IP packet can be fragmented. The default is 24 packets.
- Timeout—*Display only.* Displays the number of seconds to wait for an entire fragmented packet to arrive. The timer starts after the first fragment of a packet arrives. If all fragments of the packet do not arrive by the number of seconds displayed, all fragments of the packet that were already received will be discarded. The default is 5 seconds.
- Threshold—*Display only.* Displays the IP packet threshold, or the limit after which no new chains can be created in the reassembly module.
- Queue—*Display only.* Displays the number of IP packets waiting in the queue for reassembly.
- Assembled—*Display only.* Displays the number of IP packets successfully reassembled.
- Fail—*Display only.* Displays the number of failed reassembly attempts.
- Overflow—*Display only.* Displays the number of IP packets in the overflow queue.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Edit Fragment

Configuration > Properties > Fragment > Edit Fragment

The Edit Fragment dialog box lets you configure the IP fragment database of the selected interface.

Fields

- Interface—Displays the interface you selected in the Fragment panel. Changes made in the Edit Fragment dialog box are applied to the interface displayed.
- Size—Sets the maximum number of packets that can be in the IP reassembly database waiting for reassembly.
- Chain Length—Sets the maximum number of packets into which a full IP packet can be fragmented.
- Timeout—Sets the maximum number of seconds to wait for an entire fragmented packet to arrive. The timer starts after the first fragment of a packet arrives. If all fragments of the packet do not arrive by the number of seconds specified, all fragments of the packet that were already received will be discarded.
- Restore Defaults—Restores the factory default settings:
  - Size is 200.
  - Chain is 24 packets.
  - Timeout is 5 seconds.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Anti-Spoofing

Configuration > Properties > Anti-Spoofing

The Anti-Spoofing window lets you enable Unicast Reverse Path Forwarding on an interface. Unicast RPF guards against IP spoofing (a packet uses an incorrect source IP address to obscure its true source) by ensuring that all packets have a source IP address that matches the correct source interface according to the routing table.

Normally, the security appliance only looks at the destination address when determining where to forward the packet. Unicast RPF instructs the security appliance to also look at the source address; this is why it is called Reverse Path Forwarding. For any traffic that you want to allow through the security appliance, the security appliance routing table must include a route back to the source address. See RFC 2267 for more information.

For outside traffic, for example, the security appliance can use the default route to satisfy the Unicast RPF protection. If traffic enters from an outside interface, and the source address is not known to the routing table, the security appliance uses the default route to correctly identify the outside interface as the source interface.

If traffic enters the outside interface from an address that is known to the routing table, but is associated with the inside interface, then the security appliance drops the packet. Similarly, if traffic enters the inside interface from an unknown source address, the security appliance drops the packet because the matching route (the default route) indicates the outside interface.

Unicast RPF is implemented as follows:

- ICMP packets have no session, so each packet is checked.
- UDP and TCP have sessions, so the initial packet requires a reverse route lookup. Subsequent packets arriving during the session are checked using an existing state maintained as part of the session. Non-initial packets are checked to ensure they arrived on the same interface used by the initial packet.

Fields

- Interface—Lists the interface names.
- Anti-Spoofing Enabled—Shows whether an interface has Unicast RPF enabled, Yes or No.
- Enable—Enables Unicast RPF for the selected interface.
- Disable—Disables Unicast RPF for the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	•	—

TCP Options

Configuration > Properties > TCP Options

The TCP Options window lets you set parameters for TCP connections.

Fields

- Inbound and Outbound Reset area—Sets whether to reset denied TCP connections for inbound and outbound traffic.
  - Interface column—Shows the interface name.
  - Inbound Reset column—Shows the interface reset setting for inbound TCP traffic, Yes or No. Enabling this setting causes the security appliance to send TCP resets for all inbound TCP sessions that attempt to transit the security appliance and are denied by the security appliance based on access lists or AAA settings. Traffic between same security level interfaces is also affected. When this option is not enabled, the security appliance silently discards denied packets.
  - Outbound Reset column—Shows the interface reset setting for outbound TCP traffic, Yes or No. Enabling this setting causes the security appliance to send TCP resets for all outbound TCP sessions that attemp to transit the security appliance and are denied by the security appliance based on access lists or AAA settings. Traffic between same security level interfaces is also affected. When this option is not enabled, the security appliance silently discards denied packets.
  - Edit button—Sets the inbound and outbound reset settings for the interface.
- Send Reset Reply for Denied Outside TCP Packets check box—Enables resets for TCP packets that terminate at the least secure interface and are denied by the security appliance based on access lists or AAA settings. When this option is not enabled, the security appliance silently discards denied packets. If you enable Inbound Resets for the least secure interface (see TCP Reset Settings), then you do not also have to enable this setting; Inbound Resets handle to-the-security appliance traffic as well as through the security appliance traffic.
- Force Maximum Segment Size for TCP check box and field—Sets the maximum TCP segment size in bytes, between 48 and any maximum number. The default value is 1380 bytes. You can disable this feature by setting the bytes to 0. Both the host and the server can set the maximum segment size when they first establish a connection. If either maximum exceeds the value you set here, then the security appliance overrides the maximum and inserts the value you set. For example, if you set a maximum size of 1200 bytes, when a host requests a maximum size of 1300 bytes, then the security appliance alters the packet to request 1200 bytes.
- Force Minimum Segment Size for TCP check box and field—Overrides the maximum segment size to be no less than the number of bytes you set, between 48 and any maximum number. This feature is disabled by default (set to 0). Both the host and the server can set the maximum segment size when they first establish a connection. If either maximum is less than the value you set for the Force Minimum Segment Size for TCP Proxy field, then the security appliance overrides the maximum and inserts the “minimum” value you set (the minimum value is actually the smallest maximum allowed). For example, if you set a minimum size of 400 bytes, if a host requests a maximum value of 300 bytes, then the security appliance alters the packet to request 400 bytes.
- Force TCP Connection to Linger in TIME_WAIT State for at Least 15 Seconds check box—Forces each TCP connection to linger in a shortened TIME_WAIT state of at least 15 seconds after the final normal TCP close-down sequence. You might want to use this feature if an end host application default TCP terminating sequence is a simultaneous close. The default behavior of the security appliance is to track the shutdown sequence and release the connection after two FINs and the ACK of the last FIN segment. This quick release heuristic enables the security appliance to sustain a high connection rate, based on the most common closing sequence, known as the normal close sequence. However, in a simultaneous close, both ends of the transaction initiate the closing sequence, as opposed to the normal close sequence where one end closes and the other end acknowledges prior to initiating its own closing sequence (see RFC 793). Thus, in a simultaneous close, the quick release forces one side of the connection to linger in the CLOSING state. Having many sockets in the

CLOSING state can degrade the performance of an end host. For example, some WinSock mainframe clients are known to exhibit this behavior and degrade the performance of the mainframe server. Using this feature creates a window for the simultaneous close down sequence to complete.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

TCP Reset Settings

Configuration > Properties > TCP Options > TCP Reset Settings

This dialog box sets the inbound and outboud reset settings for an interface.

Fields

- Send Reset Reply for Denied Inbound TCP Packets check box—Sends TCP resets for all inbound TCP sessions that attempt to transit the security appliance and are denied by the security appliance based on access lists or AAA settings. Traffic between same security level interfaces is also affected. When this option is not enabled, the security appliance silently discards denied packets.

  You might want to explicitly send resets for inbound traffic if you need to reset identity request (IDENT) connections. When you send a TCP RST (reset flag in the TCP header) to the denied host, the RST stops the incoming IDENT process so that you do not have to wait for IDENT to time out. Waiting for IDENT to time out can cause traffic to slow because outside hosts keep retransmitting the SYN until the IDENT times out, so the **service resetinbound** command might improve performance.
- Send Reset Reply for Denied Outbound TCP Packets check box—Sends TCP resets for all outbound TCP sessions that attempt to transit the security appliance and are denied by the security appliance based on access lists or AAA settings. Traffic between same security level interfaces is also affected. When this option is not enabled, the security appliance silently discards denied packets. This option is enabled by default. You might want to disable outbound resets to reduce the CPU load during traffic storms, for example.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Timeouts

Configuration > Properties > Timeouts

The Timeouts window lets you set the timeout durations for use with the security appliance. 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.

Note: It is recommended that you do not change these values unless advised to do so by Customer Support.

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 (01:00:00). Setting the value of 00:00:00 means never close this connection. To close this connection immediately after all calls are cleared, a value of 1 second (00:00:01) 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 (00:05:00). 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 (00:05:00). The minimum time is 30 seconds. Uncheck the check box to return to the default value.
- 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 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 0:03:00.
- 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 0:02:00.

- 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.

- 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 timeout.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Configuring QoS

Priority Queue

Configuration > Properties > Priority Queue

The Priority Queue window shows the priority queue parameters on each configured interface. Priority queueing is disabled by default.

Fields

- Interface—Shows the name of the interface.
- Queue Limit—Shows the maximum number of packets that can be enqueued to a normal or priority queue before it drops the connection.

Note Both queues have the same limit. Packets in the priority queue are totally drained before packets in the normal priority queue are transmitted.

- Transmission Ring Limit—Specifies the depth of the priority queues. If priority queuing is not enabled, this column shows the message: “Ring Disabled.”
- Edit—Opens the Edit Priority Queue dialog box, in which you can change the priority-queue parameters.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

Add/Edit Priority Queue

Configuration > Properties > Priority Queue > Add/Edit Priority Queue

The Add/Edit Priority Queue dialog box lets you create or change the priority queue parameters for a configured interface.

The transmission ring limit is the number of either type of packets allowed into the driver before the driver pushes back to the queues sitting in front of the interface to let them buffer packets until the congestion clears. In general, you can adjust the queue-limit and transmission ring limit parameters to optimize the flow of low-latency traffic.

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 *tail drop*. To avoid having the queue fill up, you can adjust the queue-limit parameter to increase the queue buffer size.

Fields

- Interface—Identifies the selected interface. You cannot change this field.
- Queue Limit—Specifies the maximum number of packets that can be enqueued to a normal or priority queue before it drops the connection. The minimum is 0 packets, and the maximum is 250 packets.

Note Both queues have the same limit. Packets in the priority queue are totally drained before packets in the normal priority queue are transmitted.

- Enable Transmission Ring—Lets you configure the maximum number of packets allowed in the transmit queue at any given time
- Transmission Ring Limit—Specifies 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 minimum value is 3. The upper limit of the range of values for the queue-limit and tx-ring-limit commands is determined dynamically at run time. To view this limit, enter **help** or **?** on the command line. The key determinant is the memory needed to support the queues and the memory available on the device. The queues must not exceed the available memory. The theoretical maximum number of packets is 2147483647. If priority queuing is not enabled, this column shows the message: “Ring Disabled.”

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

WCCP

Configuration > Properties > WCCP

The Web Cache Communication Protocol (WCCP) feature lets you specify WCCP service groups and redirect web cache traffic. The feature transparently redirects selected types of traffic to a group of web cache engines to optimize resource usage and lower response times.

WCCP Service Groups

Configuration > Properties > WCCP >Service Groups

The Service Groups panel lets you allocate space and enable support of the specified Web Cache Communication Protocol (WCCP) service group.

Fields

- Service—Displays the service group name or service group number for WCCP support.
- Redirect List—Displays the name of the access list that controls traffic redirected to a specific service group.
- Group List—Displays the name of the access list that determines which web caches are allowed to participate in the service group.

Add or Edit WCCP Service Group

Configuration > Properties > WCCP >Service Groups

The Add or Edit Service Group dialog box lets you change the service group parameters for a configured service group.

Fields

- Service—Specifies the service group. You can specify the web cache service, or the identification number of the service.
- Web Cache—Specifies the web cache service. The maximum number of services, including those specified with a dynamic service identifier is 256.
- Dynamic Service Number—A dynamic service identifier, which means the service definition is dictated by the cache. The dynamic service number can be from 0 to 254. This is used as the name of the service group.
- Redirect List—The predefined access list that controls traffic redirected to this service group.
- Group List—The predefined access list that determines which web caches are allowed to participate in the service group.
- Password—Enter a password up to seven characters long, which is used for MD5 authentication for messages received from the service group. The password length is one to eight characters.
- Confirm Password—Reenter the password.
- Manage—Opens the access list manager.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Redirection

Configuration > Properties > WCCP >Redirection

The Redirection panel lets you enable packet redirection on the ingress of an interface using WCCP.

Fields

- Interface—Displays the interface on which WCCP redirection is enabled.
- Service Group—Displays the name of the service group configured for WCCP..

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Add or Edit WCCP Redirection

Configuration > Properties > WCCP >Add or Edit Redirection

The Redirection panel lets you enable packet redirection on the ingress of an interface using WCCP.

Fields

- Interface—Select the interface on which to enable WCCP redirection.
- Service Group—Select the service group.
- Add Service—Opens the Add/Edit WCCP Service Group dialog box.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

WCCP

Monitoring > Properties > WCCP

The Web Cache Communication Protocol (WCCP) feature lets you monitor WCCP service groups and redirect web cache traffic. The feature transparently redirects selected types of traffic to a group of web cache engines to optimize resource usage and lower response times.

WCCP Service Groups

Monitoring > Properties > WCCP >Service Groups

The Service Groups panel lets you view allocated space and display the properties of the specified Web Cache Communication Protocol (WCCP) service group.

Fields

- Service Group—Displays the service group name or service group number for WCCP support.
- Display Mode—Select the mode to display the WCCP information in the output area. The choices are Detail, View, Service, Hash, and Buckets. The Destination address and port fields and Source address and port fields correspond only to the Hash Display mode.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Redirection

Monitoring > Properties > WCCP >Redirection

The Redirection panel lets you view details about enabled packet redirection on the ingress of an interface using WCCP.

Fields

- Show Summary—Displays summarized information about the interface on which WCCP redirection is enabled.
- Show Details—Displays detailed information about the interface on which WCCP redirection is enabled.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

CHAPTER 26

VPN

The security appliance creates a virtual private network by creating a secure connection across a TCP/IP network (such as the Internet) that users see as a private connection. It can create single-user-to-LAN connections and LAN-to-LAN connections. The secure connection is called a tunnel, and the security appliance uses tunneling protocols to negotiate security parameters, create and manage tunnels, encapsulate packets, transmit or receive them through the tunnel, and unencapsulate them. The security appliance functions as a bidirectional tunnel endpoint: it can receive plain packets, encapsulate them, and send them to the other end of the tunnel, where they are unencapsulated and sent to their final destination. It can also receive encapsulated packets, unencapsulate them, and send them to their final destination.

The security appliance performs the following VPN functions:

- Establishes tunnels.
- Negotiates tunnel parameters.
- Enforces VPN policies.
- Authenticates users.
- Authorizes users for specific levels of use and access.
- Performs accounting functions.
- Assigns user addresses.
- Encrypts and decrypts data.
- Manages security keys.
- Manages data transfer across the tunnel.
- Manages data transfer inbound and outbound as a tunnel endpoint or router.

The security appliance invokes various standard protocols to accomplish these functions.

VPN Wizard

Configuration > VPN > VPN Wizard

The VPN wizard lets you configure basic LAN-to-LAN and remote access VPN connections. Use ASDM to edit and configure advanced features.

Note The VPN wizard lets you assign either preshared keys or digital certificates for authentication. However, to use certificates, you must enroll with a certification authority and configure a trustpoint prior to using the wizard. Use the ASDM Device Administration > Certificate panels and online Help to accomplish these tasks.

VPN Overview

The security appliance creates a Virtual Private Network by creating a secure connection across a TCP/IP network (such as the Internet) that users see as a private connection. It can create single-user-to-LAN connections and LAN-to-LAN connections.

The secure connection is called a tunnel, and the security appliance uses tunneling protocols to negotiate security parameters, create and manage tunnels, encapsulate packets, transmit or receive them through the tunnel, and unencapsulate them. The security appliance functions as a bidirectional tunnel endpoint: it can receive plain packets, encapsulate them, and send them to the other end of the tunnel where they are unencapsulated and sent to their final destination. It can also receive encapsulated packets, unencapsulate them, and send them to their final destination.

The security appliance performs the following functions:

- Establishes tunnels
- Negotiates tunnel parameters
- Authenticates users
- Assigns user addresses
- Encrypts and decrypts data
- Manages security keys
- Manages data transfer across the tunnel
- Manages data transfer inbound and outbound as a tunnel endpoint or router

VPN Tunnel Type

VPN Wizard > VPN Tunnel Type

Use the VPN Tunnel Type panel to select the type of VPN tunnel to define, remote access or LAN-to-LAN, and to identify the interface that connects to the remote IPSec peer.

Fields

- Site-to-Site—Click to create a LAN-to-LAN VPN configuration. Use between two IPSec security gateways, which can include security appliances, VPN concentrators, or other devices that support site-to-site IPSec connectivity. When you select this option, the VPN wizard displays a series of panels that let you to enter the attributes a site-to-site VPN requires.
- Remote Access—Click to create a configuration that achieves secure remote access for VPN clients, such as mobile users. This option lets remote users securely access centralized network resources. When you select this option, the VPN wizard displays a series of panels that let you enter the attributes a remote access VPN requires.
- VPN Tunnel Interface—Select the interface that establishes a secure tunnel with the remote IPSec peer. If the security appliance has multiple interfaces, you need to plan the VPN configuration before running this wizard, identifying the interface to use for each remote IPSec peer with which you plan to establish a secure connection.

- Enable inbound IPSec sessions to bypass interface access lists—Enable IPSec authenticated inbound sessions to always be permitted through the security appliance (that is, without a check of the interface access-list statements). Be aware that the inbound sessions bypass only the interface ACLs. Configured group-policy, user, and downloaded ACLs still apply.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Remote Site Peer

VPN Wizard > Remote Site Peer

Use the Remote Site Peer panel for the following tasks:

1. Providing the IP address of the remote IPSec peer that terminates this VPN tunnel.
2. Creating the tunnel group for the remote peer.
3. Selecting and configuring an authentication method.

Fields

- Peer IP Address—Type the IP address of the remote IPSec peer that terminates the VPN tunnel. The peer might be another security appliance, a VPN concentrator, or any other gateway device that supports IPSec.
- Authentication Method—The remote site peer authenticates either with a preshared key or a certificate.
  - Pre-shared Key—Click to use a preshared key for authentication between the local security appliance and the remote IPSec peer.

    Using a preshared key is a quick and easy way to set up communication with a limited number of remote peers and a stable network. It may cause scalability problems in a large network because each IPSec peer requires configuration information for each peer with which it establishes secure connections.

    Each pair of IPSec peers must exchange preshared keys to establish secure tunnels. Use a secure method to exchange the preshared key with the administrator of the remote site.
  - Pre-shared Key—Type the preshared key. Maximum 127 characters.
  - Certificate—Click to use certificates for authentication between the local security appliance and the remote IPSec peer. To complete this section, you must have previously enrolled with a CA and downloaded one or more certificates to the security appliance.

    Digital certificates are an efficient way to manage the security keys used to establish an IPSec tunnel. A digital certificate contains information that identifies a user or device, such as a name, serial number, company, department or IP address. A digital certificate also contains a copy of the owner's public key.

To use digital certificates, each peer enrolls with a certification authority (CA), which is responsible for issuing digital certificates. A CA can be a trusted vendor or a private CA that you establish within an organization.

When two peers want to communicate, they exchange certificates and digitally sign data to authenticate each other. When you add a new peer to the network, it enrolls with a CA, and none of the other peers require additional configuration.

- Certificate Signing Algorithm—Select the algorithm for signing digital certificates, either rsa-sig for RSA or dsa-sig for DSA.
- Trustpoint Name—Select the name that identifies the certificate the security appliance sends to the remote peer. This list displays trustpoints with a certificate of the type previously selected in the certificate signing algorithm list.
- Challenge/response authentication (CRACK)—Provides strong mutual authentication when the client authenticates using a popular method such as RADIUS and the server uses public key authentication. The security appliance supports CRACK as an IKE option in order to authenticate the Nokia VPN Client on Nokia 92xx Communicator Series devices.

- Tunnel Group Name—Type a name to create the record that contains tunnel connection policies for this IPSec connection. A tunnel group can specify authentication, authorization, and accounting servers, a default group policy, and IKE attributes. A tunnel group that you configure with this VPN wizard specifies an authentication method, and uses the security appliance Default Group Policy.

  By default, ASDM populates this box with the value of the Peer IP address. You can change this name. Maximum 64 characters.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

IKE Policy

VPN Wizard >IKE Policy

IKE, also called Internet Security Association and Key Management Protocol (ISAKMP), is the negotiation protocol that lets two hosts agree on how to build an IPSec Security Association. Each IKE negotiation is divided into two sections called Phase1 and Phase 2.

- Phase 1 creates the first tunnel, which protects later IKE negotiation messages.
- Phase 2 creates the tunnel that protects data.

Use the IKE Policy panel to set the terms of the Phase 1 IKE negotiations, which include the following:

- An encryption method to protect the data and ensure privacy.
- An authentication method to ensure the identity of the peers.
- A Diffie-Hellman group to establish the strength of the of the encryption-key-determination algorithm. The security appliance uses this algorithm to derive the encryption and hash keys.

Fields

- Encryption—Select the symmetric encryption algorithm the security appliance uses to establish the Phase 1 SA that protects Phase 2 negotiations. The security appliance supports the following encryption algorithms:

Algorithm	Explanation
DES	Data Encryption Standard. Uses a 56-bit key.
3DES	Triple DES. Performs encryption three times using a 56-bit key.
AES-128	Advanced Encryption Standard. Uses a 128-bit key.
AES-192	AES using a 192-bit key.
AES-256	AES using a 256-bit key

  The default, 3DES, is more secure than DES but requires more processing for encryption and decryption. Similarly, the AES options provide increased security, but also require increased processing.

- Authentication—Select the hash algorithm used for authentication and ensuring data integrity. The default is SHA. MD5 has a smaller digest and is considered to be slightly faster than SHA. There has been a demonstrated successful (but extremely difficult) attack against MD5. However, the Keyed-Hash Message Authentication Code (HMAC) version used by the security appliance prevents this attack.
- DH Group—Select the Diffie-Hellman group identifier, which the two IPSec peers use to derive a shared secret without transmitting it to each other. The default, Group 2 (1024-bit Diffie-Hellman), requires less CPU time to execute but is less secure than Group 5 (1536-bit). Group 7 is for use with the Movian VPN client, but works with any peer that supports Group 7 (ECC).

Note The default value for the VPN 3000 Series Concentrator is MD5. A connection between the security appliance and the VPN Concentrator requires that the authentication method for Phase I and II IKE negotiations be the same on both sides of the connection.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

IPSec Encryption and Authentication

VPN Wizard > IPSec Encryption and Authentication

Use this IPSec Encryption and Authentication panel to select the encryption and authentication methods to use for Phase 2 IKE negotiations, which create the secure VPN tunnel. These values must be exactly the same for both peers.

Fields

- Encryption—Select the symmetric encryption algorithm the security appliance uses to establish the VPN tunnel. The security appliance uses encryption to protect the data that travels across the tunnel and ensure privacy. Valid encryption methods include the following:

Encryption Method	Explanation
DES	Data Encryption Standard. Uses a 56-bit key.
3DES	Triple DES. Encrypts three times using a 56-bit key.
AES-128	Advanced Encryption Standard. Uses a 128-bit key.
AES-192	AES using a 192-bit key.
AES-256	AES using a 256-bit key

  The default, 3DES, is more secure than DES but requires more processing for encryption and decryption. Similarly, the AES options provide increased security, but also require increased processing.

- Authentication—Select the hash algorithm used for authentication and ensuring data integrity. The default is SHA. MD5 has a smaller digest and is considered to be slightly faster than SHA. There has been a demonstrated successful (but extremely difficult) attack against MD5. However, the Keyed-Hash Message Authentication Code (HMAC) version used by the security appliance prevents this attack.

Note The default value for the VPN 3000 Series Concentrator is MD5. A connection between the security appliance and the VPN Concentrator requires that the authentication method for Phase I and Phase II IKE negotiations be the same on both sides of the connection.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Local Hosts and Networks

VPN Wizard > Hosts and Networks

Use the Hosts and Networks panel to identify local and remote hosts and networks that can use this LAN-to-LAN IPSec tunnel to send and receive data. You can identify hosts and networks by IP address, DNS name or group policy. Depending on that choice, the remaining fields in this panel change.

For IPSec to succeed, both peers in the LAN-to-LAN connection must have compatible entries for hosts and networks. The hosts and networks you configure as Local Hosts and Networks in this panel must be configured as Remote Hosts and Networks on the device at the remote site for the LAN-to-LAN connection. The local security appliance and the remote device must have at least one transform set in common for this LAN-to-LAN connection.

Fields

- Source area—Lets you configure the local hosts and networks.
  - Type—Select either any, IP Address, Network Object Group, or Interface IP.

    Selecting IP Address displays IP Address and Netmask fields. Enter the IP address of the host or network. Either type the IP address or click the adjacent ... button to select a host or network. Select the subnet mask for the IP address. If you used the ... button to select a host or network, ASDM completes this box automatically.

    Selecting Network Object Group displays the Group Name field, where you specify a group name. This option lets you configure an entire group to use the tunnel. You configure these groups in the Configuration > Features > Building Blocks > Hosts and Networks panel.

    Selecting Interface IP displays the Interface field, where you select the interface name.
- Destination area—Lets you configure the local hosts and networks.
  - Type—Select either any, IP Address, Network Object Group, or Interface IP.

    Selecting IP Address displays IP Address and Netmask fields. Enter the IP address of the host or network. Either type the IP address or click the adjacent ... button to select a host or network. Select the subnet mask for the IP address. If you used the ... button to select a host or network, ASDM completes this box automatically.

    Selecting Network Object Group displays the Group Name field, where you specify a group name. This option lets you configure an entire group to use the tunnel. You configure these groups in the Configuration > Features > Building Blocks > Hosts and Networks panel.

    Selecting Interface IP displays the Interface field, where you select the interface name.
- Exempt ASA side host/network from address translation—Allows traffic to pass through the security appliance without address translation.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Summary

VPN Wizard > Summary

The Summary panel displays all of the attributes of this VPN LAN-to-LAN connection as configured.

Fields

Back—To make changes, click **Back** until you reach the appropriate panel.

Finish—When you are satisfied with the configuration, click **Finish**. ASDM saves the LAN-to-LAN configuration. After you click **Finish**, you can no longer use the VPN wizard to make changes to this configuration. Use ASDM to edit and configure advanced features.

Cancel—To remove the configuration, click **Cancel**.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Remote Access Client

VPN Wizard > Remote Access Client

Use the Remote Access Client panel to identify the type of remote access users this connection serves.

Fields

- Cisco VPN Client Release 3.x or higher, or other Easy VPN Remote product—Click for IPSec connections, including compatible software and hardware clients other than those named here.
- Microsoft Windows client using L2TP over IPSec—Click to enable connections from Microsoft Windows and Microsoft Windows Mobile clients over a public IP network. L2TP uses PPP over UDP (port 1701) to tunnel the data. Enable one or more of the following PPP authentication protocols:
  - PAP—Passes cleartext username and password during authentication and is not secure.
  - CHAP—In response to the server challenge, the client returns the encrypted [challenge plus password] with a cleartext username. This protocol is more secure than the PAP, but it does not encrypt data.
  - MS-CHAP, Version 1—Similar to CHAP but more secure in that the server stores and compares only encrypted passwords rather than cleartext passwords as in CHAP.
  - MS-CHAP, Version 2—Contains security enhancements over MS-CHAP, Version 1.
  - EAP—Enables EAP which permits the security appliance to proxy the PPP authentication process to an external RADIUS authentication server.
- Client will send tunnel group name as username@tunnelgroup—Check to enable the security appliance to associate different users that are establishing L2TP over IPSec connections with different tunnel groups. Since each tunnel group has its own AAA server group and IP address pools, users can be authenticated through methods specific to their tunnel group.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

VPN Client Authentication Method and Tunnel Group Name

VPN Wizard > VPN Client Authentication Method and Tunnel Group Name

Use the VPN Client Authentication Method and Tunnel Group Name panel to configure an authentication method and create a tunnel group.

Fields

- Authentication Method—The remote site peer authenticates either with a preshared key or a certificate.
  - Pre-shared Key—Click to use a preshared key for authentication between the local security appliance and the remote IPSec peer.

    Using a preshared key is a quick and easy way to set up communication with a limited number of remote peers and a stable network. It may cause scalability problems in a large network because each IPSec peer requires configuration information for each peer with which it establishes secure connections.

    Each pair of IPSec peers must exchange preshared keys to establish secure tunnels. Use a secure method to exchange the preshared key with the administrator of the remote site.
  - Pre-shared Key—Type the preshared key.
  - Certificate—Click to use certificates for authentication between the local security appliance and the remote IPSec peer. To complete this section, you must have previously enrolled with a CA and downloaded one or more certificates to the security appliance.

    Digital certificates are an efficient way to manage the security keys used to establish an IPSec tunnel. A digital certificate contains information that identifies a user or device, such as a name, serial number, company, department or IP address. A digital certificate also contains a copy of the owner's public key.

    To use digital certificates, each peer enrolls with a certification authority (CA), which is responsible for issuing digital certificates. A CA can be a trusted vendor or a private CA that you establish within an organization.

    When two peers want to communicate, they exchange certificates and digitally sign data to authenticate each other. When you add a new peer to the network, it enrolls with a CA, and none of the other peers require additional configuration.
  - Trustpoint Name—Select the name that identifies the certificate the security appliance sends to the remote peer.
  - Certificate Signing Algorithm—Select the algorithm for signing digital certificates, either rsa-sig for RSA or dsa-sig for DSA.
  - Challenge/response authentication (CRACK)—Provides strong mutual authentication when the client authenticates using a popular method such as RADIUS and the server uses public key authentication. The security appliance supports CRACK as an IKE option in order to authenticate the Nokia VPN Client on Nokia 92xx Communicator Series devices.

- Tunnel Group Name—Type a name to create the record that contains tunnel connection policies for this IPSec connection. A tunnel group can specify authentication, authorization, and accounting servers, a default group policy, and IKE attributes. A tunnel group that you configure with this VPN wizard specifies an authentication method, and uses the security appliance Default Group Policy.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Client Authentication

VPN Wizard > Client Authentication

Use the Client Authentication panel to select the method by which the security appliance authenticates remote users.

Fields

Select one of the following options:

- Authenticate using the local user database—Click to use authentication internal to the security appliance. Use this method for environments with a small, stable number of users. The next panel lets you create accounts on the security appliance for individual users.
- Authenticate using an AAA server group—Click to use an external server group for remote user authentication.
- AAA Server Group—Select a AAA server group configured previously.
- New ...—Click to configure a new AAA server group.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

New Authentication Server Group

VPN Wizard > New Authentication Server Group

User the New Authentication Server Group panel to define one or more new AAA servers.

Fields

To configure a new AAA server group that contains just one server, provide the following information:

- Server Group Name—Type a name for the server group. You associate this name with users whom you want to authenticate using this server.
- Authentication Protocol—Select the authentication protocol the server uses. Options include TACACS+, RADIUS, SDI, NT, and Kerberos.
- Server IP Address—Type the IP address for the AAA server.
- Interface—Select the security appliance interface on which the AAA server resides.
- Server Secret Key—Type a case-sensitive, alphanumeric keyword of up to 127 characters. The server and security appliance use the key to encrypt data that travels between them. The key must be the same on both the security appliance and server. You can use special characters, but not spaces.
- Confirm Server Secret Key—Type the secret key again.

To add more servers to this new group, or to change other AAA server settings, go to Configuration > Features > Properties > AAA.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

User Accounts

VPN Wizard > User Accounts

Use the User Accounts panel to add new users to the security appliance internal user database for authentication purposes.

Fields

Provide the following information:

- User to Be Added—Use the fields in this section to add a user.
  - Username—Enter the username.
  - Password—(Optional) Enter a password.
  - Confirm Password—(Optional) Reenter the password.
- Add — Click to add a user to the database after you have entered the username and optional password.
- Username—Displays the names of all users in the database.
- Delete—To remove a user from the database, highlight the appropriate username and click **Delete**.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Address Pool

VPN Wizard > Address Pool

Use the Address Pool panel to configure a pool of local IP addresses that the security appliance assigns to remote VPN clients.

Fields

- Tunnel Group Name—Displays the name of the tunnel group to which the address pool applies. You set this name in the VPN Client Tunnel Group Name and Authentication Method panel.
- Pool Name—Select a descriptive identifier for the address pool. The security appliance associates the pool name with a tunnel group.
- Range Start Address—Type the starting IP address in the address pool.
- Range End Address—Type the ending IP address in the address pool.
- Subnet Mask—(Optional) Select the subnet mask for these IP addresses

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Attributes Pushed to Client

VPN Wizard > Attributes Pushed to Client

Use the Attributes Pushed to Client **(Optional)** panel to have the security appliance pass information about DNS and WINS servers and the default domain name to remote access clients.

Fields

Provide information for remote access clients to use.

- Tunnel Group—Displays the name of the tunnel group to which the address pool applies. You set this name in the VPN Client Tunnel Group Name and Authentication Method panel.
- Primary DNS Server—Type the IP address of the primary DNS server.
- Secondary DNS Server—Type the IP address of the secondary DNS server.
- Primary WINS Server—Type the IP address of the primary WINS server.

- Secondary WINS Server— Type the IP address of the secondary WINS server.
- Default Domain Name—Type the default domain name. Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Address Translation Exemption

VPN Wizard > Address Translation Exemption

Use the Address Translation Exemption (Optional) panel to identify local hosts/networks which do not require address translation. By default, the security appliance hides the real IP addresses of internal hosts and networks from outside hosts by using dynamic or static Network Address Translation (NAT). NAT minimizes risks of attack by untrusted outside hosts, but may be improper for those who have been authenticated and protected by VPN.

For example, an inside host using dynamic NAT has its IP address translated by matching it to a randomly selected address from a pool. Only the translated address is visible to the outside. Remote VPN clients that attempt to reach these hosts by sending data to their real IP addresses cannot connect to these hosts, unless you configure a NAT exemption rule.

Note If you want all hosts and networks to be exempt from NAT, configure nothing on this panel. If you have even one entry, all other hosts and networks are subject to NAT.

Fields

- Host/Network to Be Added—Complete these fields to exempt a particular host or network from NAT.
  - IP Address—Click to identify hosts and networks by IP address.
  - Name—Click to identify hosts by their hostnames.
  - Group—Click to identify hosts and networks by tunnel group. This option lets you configure an entire group to use the tunnel.
  - Group—Select the name of the tunnel group.
  - Interface—Select the name of the interface that connects to the hosts or networks you have selected.
  - IP address—Select the IP address of the host or network. Either type the IP address or click the adjacent ... button to view a diagram of the network and select a host or network.
  - Mask—Select the subnet mask for the IP address.
  - Name—Select the hostname. Use fully qualified domain names.
- Add—Click to add the host or network the Selected Hosts/Networks list after you have completed the applicable fields.

- Selected Hosts/Networks—Displays the hosts and networks that are exempt from NAT. If you want all hosts and networks to be exempt from NAT, leave this list empty.
- Enable split tunneling—Select to have traffic from remote access clients destined for the public Internet sent unencrypted. Split tunneling causes traffic for protected networks to be encrypted, while traffic to unprotected networks is unencrypted. When you enable split tunneling, the security appliance pushes a list of IP addresses to the remote VPN client after authentication. The remote VPN client encrypts traffic to the IP addresses that are behind the security appliance. All other traffic travels unencrypted directly to the Internet without involving the security appliance.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

IKE

IKE, also called ISAKMP, is the negotiation protocol that lets two hosts agree on how to build an IPSec security association. To configure the security appliance for virtual private networks, you set global IKE parameters that apply system wide, and you also create IKE policies that the peers negotiate to establish a VPN tunnel.

Certificate Group Matching

Certificate Group Matching lets you define rules to match a user's certificate to a permission group based on fields in the distinguished name (DN). You can use any field of the certificate or you can have all certificate users share a permission group.

To match user permission groups based on fields of the certificate, you must define rules that specify the fields to match for a group and then enable each rule for that selected group. A group must already exist in the configuration before you can create a rule for it. If the group does not exist in the configuration, you must define it by using **Configuration > VPN > General > Tunnel Group**.

Once you have defined rules, you must configure a certificate group matching policy to define the method to use for identifying the permission groups of certificate users: match the group from the rules, match the group from the OU field, or use a default group for all certificate users. You can use any or all of these methods.

Policy

Configuration > VPN > IKE > Certificate Group Matching > Policy

A certificate group matching policy defines the method to use for identifying the permission groups of certificate users. You can use any or all of these methods.

Fields

- **Use the configured rules to match a certificate to a group**—Lets you use the rules you have defined under **Rules**.
- **Use the certificate OU field to determine the group**—Lets you use the organizational unit field to determine the group to which to match the certificate. This is selected by default.
- **Use the IKE identity to determine the group**—Lets you use the identity you previously defined under **Configuration > VPN > IKE > Global Parameters**. The IKE identity can be hostname, IP address, key ID, or automatic.

- **Use the peer IP address to determine the group**—Lets you use the peer's IP address. This is selected by default.
- **Default to group—**Lets you select a default group for certificate users that is used when none of the preceding methods resulted in a match. This is selected by default. Click the default group in the **Default to group** list. The group must already exist in the configuration. If the group does not appear in the list, you must define it by using **Configuration > VPN > General > Tunnel Group**.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Rules

Configuration > VPN > IKE > Certificate Group Matching > Rules

The **Certificate Group Matching Rules** panel lets you add, edit, and delete rules to/from your configuration.

To match user permission groups based on fields of the certificate, you must define rules that specify which fields to match for a group and then enable each rule for that selected group. Rules cannot be longer than 255 characters. A group must already exist in the configuration before you can create a rule for it.

You can assign multiple rules to the same group. To do that, you add the rule priority and group first. Then you define as many criteria statements as you need for each group. When multiple rules are assigned to the same group, a match results for the first rule that tests true.

To match user permission groups based on multiple fields in the certificate so that all the criteria must match for the user to be assigned to a permission group, create a single rule with multiple matching criteria. To match user permission groups based on one criterion or another so that successfully matching any of the criteria identifies the member of the group, create multiple rules.

Fields

- **Map Name**—Displays the names of configured certificate-to-group maps.
- **Rule Priority**—Displays the importance of the rule as a matching criterion. The lower the value, the higher the priority. The default value is 10.
- **Mapped to Group**—Displays the group to which the rule is mapped.
- **Field—**Displays the type of distinguished name (Subject or Issuer) used in the rule.
- **Component**—Displays the distinguished name component used in the rule. For a list of possibilities and their definitions, see **Add Certificate Matching Rule Criterion Help**.
- **Operator**—Displays the operator used in the rule: Equals, Not Equals, Contains, and Does Not Contain.
- **Value**—Displays the value to be matched against.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

Add/Edit Certificate Matching Rule

Configuration > VPN > IKE > Certificate Group Matching > Rules > Add/Edit Certificate Matching Rule

Use the **Add/Edit Certificate Matching Rule** dialog box to define a certificate matching rule.

Fields

- **Map Existing**—Select the name of the map to include the rule.
- **Map New**—Enter a new map name for a rule.
- **Rule Priority**—Specify a number that indicates the level of priority for this rule. For the first rule defined, the default priority is 10. Each rule must have a unique priority. The lower the value, the higher the priority.
- **Mapped to Group**—Select the tunnel group to map to this rule. Groups must already exist in the configuration. If the group does not appear in the list, you must define it by using **Configuration > VPN > General > Tunnel Group**.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	—	—

Add/Edit Certificate Matching Rule Criterion

Configuration > VPN > IKE > Certificate Group Matching > Rules > Add/Edit Certificate Matching Rule Criterion

Use the **Add/Edit Certificate Matching Rule Criterion** dialog box to configure a certificate matching rule criterion for the selected group.

Fields

- **Field**—Specify the type of distinguished name to be used in the rule: Subject and Issuer, which consist of a specific-to-general identification hierarchy: CN, OU, O, L, SP, and C. These labels and acronyms conform to X.500 terminology.

  - **Subject**—The person or system that uses the certificate. For a CA root certificate, the Subject and Issuer are the same.
  - **Issuer**—The CA or other entity (jurisdiction) that issued the certificate.
- **Component**—Select the distinguished name component used in the rule:

DN Field	Definition
Whole Field	The entire DN.
Country (C)	The two-letter country abbreviation. These codes conform to ISO 3166 country abbreviations.
Common Name (CN)	The name of a person, system, or other entity. This is the lowest (most specific) level in the identification hierarchy.
DN Qualifier (DNQ)	A specific DN attribute.
E-mail Address (EA)	The e-mail address of the person, system or entity that owns the certificate.
Generational Qualifier (GENQ)	A generational qualifier such as Jr., Sr., or III.
Given Name (GN)	The first name of the certificate owner.
Initials (I)	The first letters of each part of the certificate owner's name.
Locality (L)	The city or town where the organization is located.
Name (N)	The name of the certificate owner.
Organization (O)	The name of the company, institution, agency, association, or other entity.
Organizational Unit (OU)	The subgroup within the organization.
Serial Number (SER)	The serial number of the certificate.
Surname (SN)	The family name or last name of the certificate owner.
State/Province (S/P)	The state or province where the organization is located.
Title (T)	The title of the certificate owner, such as Dr.
User ID (UID)	The identification number of the certificate owner.

- **Operator**—Select the operator used in the rule:
  - **Equals**—The distinguished name field must exactly match the value.
  - **Contains**—The distinguished name field must include the value within it.
  - **Does Not Equal**—The distinguished name field must not match the value
  - **Does Not Contain**—The distinguished name field must not include the value within it.
- **Value**—Specify the value to match against.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Global Parameters

Configuration > VPN > IKE > Global Parameters

This panel lets you set system wide values for VPN tunnels. The following sections describe each of the options.

Enabling IKE on Interfaces

You must enable IKE for each interface that you want to use for VPN tunnels.

Enabling IPSec over NAT-T

NAT-T lets IPSec peers establish both remote access and LAN-to-LAN connections through a NAT device. It does this by encapsulating IPSec traffic in UDP datagrams, using port 4500, thereby providing NAT devices with port information. NAT-T auto-detects any NAT devices, and only encapsulates IPSec traffic when necessary. This feature is disabled by default.

- The security appliance can simultaneously support standard IPSec, IPSec over TCP, NAT-T, and IPSec over UDP, depending on the client with which it is exchanging data.
- When both NAT-T and IPSec over UDP are enabled, NAT-T takes precedence.
- When enabled, IPSec over TCP takes precedence over all other connection methods.

The security appliance implementation of NAT-T supports IPSec peers behind a single NAT/PAT device as follows:

- One LAN-to-LAN connection.
- Either a LAN-to-LAN connection or multiple remote access clients, but not a mixture of both.

To use NAT-T you must:

- Open port 4500 on the security appliance.
- Enable IPSec over NAT-T globally in this panel.
- Select the second or third option for the Fragmentation Policy parameter in the **Configuration > VPN > IPSec > Pre-Fragmentation** panel. These options let traffic travel across NAT devices that do not support IP fragmentation; they do not impede the operation of NAT devices that do support IP fragmentation.

Enabling IPSec over TCP

IPSec over TCP enables a VPN client to operate in an environment in which standard ESP or IKE cannot function, or can function only with modification to existing firewall rules. IPSec over TCP encapsulates both the IKE and IPSec protocols within a TCP packet, and enables secure tunneling through both NAT and PAT devices and firewalls. This feature is disabled by default.

Note This feature does not work with proxy-based firewalls.

IPSec over TCP works with remote access clients. It works on all physical and VLAN interfaces. It is a client to security appliance feature only. It does not work for LAN-to-LAN connections.

- The security appliance can simultaneously support standard IPSec, IPSec over TCP, NAT-Traversal, and IPSec over UDP, depending on the client with which it is exchanging data.
- The VPN 3002 hardware client, which supports one tunnel at a time, can connect using standard IPSec, IPSec over TCP, NAT-Traversal, or IPSec over UDP.
- When enabled, IPSec over TCP takes precedence over all other connection methods.

You enable IPSec over TCP on both the security appliance and the client to which it connects.

You can enable IPSec over TCP for up to 10 ports that you specify. If you enter a well-known port, for example port 80 (HTTP) or port 443 (HTTPS), the system displays a warning that the protocol associated with that port will no longer work. The consequence is that you can no longer use a browser to manage the security appliance through the IKE-enabled interface. To solve this problem, reconfigure the HTTP/HTTPS management to different ports.

You must configure TCP port(s) on the client as well as on the security appliance. The client configuration must include at least one of the ports you set for the security appliance.

Determining ID Method

During IKE negotiations the peers must identify themselves to each other. You can choose the identification methods from the following options:

Address	Uses the IP addresses of the hosts exchanging ISAKMP identity information.
Hostname	Uses the fully-qualified domain name of the hosts exchanging ISAKMP identity information (default). This name comprises the hostname and the domain name.
Key ID	Uses the string the remote peer uses to look up the preshared key.
Automatic	Determines IKE negotiation by connection type: • IP address for preshared key • Cert DN for certificate authentication.

Disabling Inbound Aggressive Mode Connections

Phase 1 IKE negotiations can use either Main mode or Aggressive mode. Both provide the same services, but Aggressive mode requires only two exchanges between the peers, rather than three. Aggressive mode is faster, but does not provide identity protection for the communicating parties. It is therefore necessary that they exchange identification information prior to establishing a secure SA in which to encrypt in formation. This feature is disabled by default.

Alerting Peers Before Disconnecting

Client or LAN-to-LAN sessions may be dropped for several reasons, such as: a security appliance shutdown or reboot, session idle timeout, maximum connection time exceeded, or administrator cut-off.

The security appliance can notify qualified peers (in LAN-to-LAN configurations), VPN Clients and VPN 3002 Hardware Clients of sessions that are about to be disconnected, and it conveys to them the reason. The peer or client receiving the alert decodes the reason and displays it in the event log or in a pop-up panel. This feature is disabled by default.

This panel lets you enable the feature so that the security appliance sends these alerts, and conveys the reason for the disconnect.

Qualified clients and peers include the following:

- Security appliance devices with Alerts enabled.

- VPN clients running 4.0 or later software (no configuration required).
- VPN 3002 hardware clients running 4.0 or later software, and with Alerts enabled.
- VPN 3000 Series Concentrators running 4.0 or later software, with Alerts enabled.

Waiting for Active Sessions to Terminate Prior to Reboot

You can schedule a security appliance reboot to occur only when all active sessions have terminated voluntarily. This feature is disabled by default.

Fields

- **Enable IKE**—Shows IKE status for all configured interfaces.
  - **Interface**—Displays names of all configured security appliance interfaces.
  - **IKE Enabled**—Shows whether IKE is enabled for each configured interface.
  - **Enable/Disable**s—Click to enable or disable IKE for the highlighted interface.
- **NAT Transparency**—Lets you enable or disable IPSec over NAT-T and IPSec over TCP.
  - **Enable IPSec over NAT-T**—Select to enable IPSec over NAT-T.
  - **NAT Keepalive**—Type the number of seconds that can elapse with no traffic before the security appliance terminates the NAT-T session. The default is 20 seconds. The range is 10 to 3600 seconds (one hour).
  - **Enable IPSec over TCP**—Select to enable IPSec over TCP.
  - **Enter up to 10 comma-separated TCP port values**—Type up to 10 ports on which to enable IPSec over TCP. Use a comma to separate the ports. You do not need to use spaces. The default port is 10,000. The range is 1 to 65,635.
- **Identity to Be Sent to Peer**—Lets you set the way that IPSec peers identify themselves to each other.
  - **Identity**—Select one of the following methods by which IPSec peers identify themselves:

Address	Uses the IP addresses of the hosts.
Hostname	Uses the fully-qualified domain names of the hosts. This name comprises the hostname and the domain name.
Key ID	Uses the string the remote peer uses to look up the preshared key.
Automatic	Determines IKE negotiation by connection type: IP address for preshared key or cert DN for certificate authentication.

  - **Key Id String**—Type the alpha-numeric string the peers use to look up the preshared key.
- **Disable inbound aggressive mode connections**—Select to disable aggressive mode connections.
- **Alert peers before disconnecting**—Select to have the security appliance notify qualified LAN-to-LAN peers and remote access clients before disconnecting sessions.
- **Wait for all active sessions to voluntarily terminate before rebooting**—Select to have the security appliance postpone a scheduled reboot until all active sessions terminate.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Policies

Configuration > VPN > IKE > Policies

Each IKE negotiation is divided into two sections called Phase1 and Phase 2.

Phase 1 creates the first tunnel, which protects later IKE negotiation messages. Phase 2 creates the tunnel that protects data.

To set the terms of the IKE negotiations, you create one or more IKE policies, which include the following:

- A unique priority (1 through 65,543, with 1 the highest priority).
- An authentication method, to ensure the identity of the peers.
- An encryption method, to protect the data and ensure privacy.
- An HMAC method to ensure the identity of the sender, and to ensure that the message has not been modified in transit.
- A Diffie-Hellman group to establish the strength of the of the encryption-key-determination algorithm. The security appliance uses this algorithm to derive the encryption and hash keys.
- A limit for how long the security appliance uses an encryption key before replacing it.

If you do not configure any IKE policies, the security appliance uses the default policy, which is always set to the lowest priority, and which contains the e default value for each parameter. If you do not specify a value for a specific parameter, the default value takes effect.

When IKE negotiation begins, the peer that initiates the negotiation sends all of its policies to the remote peer, and the remote peer searches for a match with its own policies, in priority order.

A match between IKE policies exists if they have the same encryption, hash, authentication, and Diffie-Hellman values, and an SA lifetime less than or equal to the lifetime in the policy sent. If the lifetimes are not identical, the shorter lifetime—from the remote peer policy—applies. If no match exists, IKE refuses negotiation and the IKE SA is not established.

Fields

- **Policies**—Displays parameter settings for each configured IKE policy.
  - **Priority #**—Shows the priority of the policy.
  - **Encryption**—Shows the encryption method.
  - **Hash**—Shows the has algorithm.
  - **D-H Group**—Shows the Diffie-Hellman group.
  - **Authentication**—Shows the authentication method.
  - **Lifetime** (secs)—Shows the SA lifetime in seconds.
- **Add/Edit/Delete**—Click to add, edit, or delete an IKE policy.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit IKE Policy

Configuration > VPN > IKE > Policies > Add/Edit IKE Policy

Fields

Priority #—Type a number to set a priority for the IKE policy. The range is 1 to 65,543, with 1 the highest priority.

Encryption—Select an encryption method. This is a symmetric encryption method that protects data transmitted between two IPSec peers.The choices follow:

des	56-bit DES-CBC. Less secure but faster than the alternatives. The default.
3des	168-bit Triple DES.
aes	128-bit AES.
aes-192	192-bit AES.
aes-256	256-bit AES.

Hash—Select the hash algorithm that ensures data integrity. It ensures that a packet comes from whom you think it comes from, and that it has not been modified in transit.

sha	SHA-1	The default is SHA-1. MD5 has a smaller digest and is considered to be slightly faster than SHA-1. A successful (but extremely difficult) attack against MD5 has occurred; however, the HMAC variant IKE uses prevents this attack.
md5	MD5	

Authentication—Select the authentication method the security appliance uses to establish the identity of each IPSec peer. Pre-shared keys do not scale well with a growing network but are easier to set up in a small network. The choices follow:

pre-share	Pre-shared keys.
rsa-sig	A digital certificate with keys generated by the RSA signatures algorithm.
crack	IKE Challenge/Response for Authenticated Cryptographic Keys protocol for mobile IPSec-enabled clients which use authentication techniques other than certificates.

D-H Group—Select the Diffie-Hellman group identifier, which the two IPSec peers use to derive a shared secret without transmitting it to each other.

1	Group 1 (768-bit)	The default, Group 2 (1024-bit Diffie-Hellman) requires less CPU time to execute but is less secure than Group 2 or 5.
2	Group 2 (1024-bit	
5	Group 5 (1536-bit)	
7	Group 7 (Elliptical curve field size is 163 bits.)	Group 7 is for use with the Movian VPN client, but with any peer that supports Group 7 (ECC).

Lifetime (secs)—Either select Unlimited or type an integer for the SA lifetime. The default is 86,400 seconds or 24 hours. With longer lifetimes, the security appliance sets up future IPSec security associations more quickly. Encryption strength is great enough to ensure security without using very fast rekey times, on the order of every few minutes. We recommend that you accept the default.

Time Measure—Select a time measure. The security appliance accepts the following values:.

120 - 86,400 seconds

2 - 1440 minutes

1 - 24 hours

1 day

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

IP Address Management

IP addresses make internetwork connections possible. They are like telephone numbers: both the sender and receiver must have an assigned number to connect. But with VPNs, there are actually two sets of addresses: the first set connects client and server on the public network; and once that connection is made, the second set connects client and server through the VPN tunnel.

In security appliance address management, we are dealing with the second set of IP addresses: those private IP addresses that connect a client with a resource on the private network, through the tunnel, and let the client function as if it were directly connected to the private network. Furthermore, we are dealing only with the private IP addresses that get assigned to clients. The IP addresses assigned to other resources on your private network are part of your network administration responsibilities, not part of security appliance management.

Therefore, when we discuss IP addresses here, we mean those IP addresses available in your private network addressing scheme, that let the client function as a tunnel endpoint.

Assignment

Configuration > VPN > IP Address Management > Assignment

The Assignment panel lets you choose a way to assign IP addresses to remote access clients.

Fields

- **Use authentication server**—Select to assign IP addresses retrieved from an authentication server on a per-user basis. If you are using an authentication server (external or internal) that has IP addresses configured, we recommend using this method. Configure AAA servers on the **Configuration > Properties > AAA Setup > AAA Servers** and **AAA Server Group** panels**.**
- **Use DHCP**— Select to obtain IP addresses from a DHCP server. If you use DHCP, configure the server on the **Configuration > Properties > DHCP Services > DHCP Server** panel.
- **Use internal address pools**—Select to have the security appliance assign IP addresses from an internally configured pool. Internally configured address pools are the easiest method of address pool assignment to configure. If you use this method, configure the IP address pools on **Configuration > VPN > IP Address Management > IP Pools** panel.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

IP Pools

Configuration > VPN > IP Address Management > IP Pools

The IP Pool box shows each configured address pool by name, and with their IP address range, for example: 10.10.147.100 to 10.10.147.177. If no pools exist, the box is empty. The security appliance uses these pools in the order listed: if all addresses in the first pool have been assigned, it uses the next pool, and so on.

If you assign addresses from a non-local subnet, we suggest that you add pools that fall on subnet boundaries to make adding routes for these networks easier.

Fields

- **Pool Name**—Displays the name of each configured address pool.
- **Starting Address**—Shows first IP address available in each configured pool.
- **Ending Address**—Shows the last IP address available in each configured pool.
- **Subnet Mask**—Shows the subnet mask for addresses in each configured pool.
- **Add**—Click to add a new address pool.
- **Edit/Delete**—Click to edit or delete an already configured address pool.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit IP Pool

Configuration > VPN > IP Address Management > IP Pools > Add/Edit IP Pool

These panels let you:

- Add a new pool of IP addresses from which the security appliance assigns addresses to clients.
- Modify an IP address pool that you have previously configured.

The IP addresses in the pool range must not be assigned to other network resources.

Fields

- **Name**—Assign an alpha-numeric name to the address pool. Limit 64 characters
- **Starting IP Address**—Enter the first IP address available in this pool. Use dotted decimal notation, for example: 10.10.147.100.
- **Ending IP Address**—Enter the last IP address available in this pool. Use dotted decimal notation, for example: 10.10.147.100.
- **Subnet Mas**k—Select the subnet mask for the IP address pool.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

IPSec

IPSec provides the most complete architecture for VPN tunnels, and it is perceived as the most secure protocol. Both LAN-to-LAN connections and client-to-LAN connections can use IPSec.

In IPSec terminology, a “peer” is a remote-access client or another secure gateway. During tunnel establishment with IPSec, the two peers negotiate security associations that govern authentication, encryption, encapsulation, and key management. These negotiations involve two phases: first, to establish the tunnel (the IKE SA); and second, to govern traffic within the tunnel (the IPSec SA).

In IPSec LAN-to-LAN connections, the security appliance can function as initiator or responder. In IPSec client-to-LAN connections, the security appliance functions only as responder. Initiators propose SAs; responders accept, reject, or make counter-proposals—all in accordance with configured SA parameters. To establish a connection, both entities must agree on the SAs.

The VPN Client complies with the IPSec protocol and is specifically designed to work with the security appliance. However, the security appliance can establish IPSec connections with many protocol-compliant clients. Likewise, the security appliance can establish LAN-to-LAN connections with other protocol-compliant VPN devices, often called secure gateways.

Thesecurity appliance supports these IPSec attributes:

- Main mode for negotiating phase one ISAKMP security associations when using digital certificates for authentication
- Aggressive mode for negotiating phase one ISAKMP Security Associations (SAs) when using preshared keys for authentication
- Authentication Algorithms:
  - ESP-MD5-HMAC-128
  - ESP-SHA1-HMAC-160
- Authentication Modes:
  - Preshared Keys
  - X.509 Digital Certificates
- Diffie-Hellman Groups 1, 2, 5, and 7
- Encryption Algorithms:
  - AES-128, -192, and -256
  - 3DES-168
  - DES-56
  - ESP-NULL
- Extended Authentication (XAuth)
- Mode Configuration (also known as ISAKMP Configuration Method)
- Tunnel Encapsulation Mode
- IP compression (IPCOMP) using LZS

IPSec Rules

Configuration > VPN > IPSec > IPSec Rules

This pane shows the currently configured IPSec rules. Use it to add, edit, delete and move up, move down, cut, copy, and paste an IPSec rule.

Fields

Note You cannot edit, delete, or copy an implicit rule. The security appliance implicitly accepts the traffic selection proposal from remote clients when configured with a dynamic tunnel policy. You can override it by giving a specific traffic selection.

- **Type: Priority**—Displays the type of rule (static or dynamic) and its priority.
- **Traffic Selection**
  - **#**—Indicates the rule number.
  - **Source**—Indicates the IP addresses that are subject to this rule when traffic is sent to the IP addresses listed in the **Remote Side Host/Network** column. In detail mode (see the **Show Detail** button), an address column might contain an interface name with the word **any**, such as **inside:any**. **any** means that any host on the inside interface is affected by the rule.
  - **Destination**—Lists the IP addresses that are subject to this rule when traffic is sent from the IP addresses listed in the **Security Appliance Side Host/Network** column. In detail mode (see the **Show Detail** button), an address column might contain an interface name with the word **any**, such as **outside:any**. **any** means that any host on the outside interface is affected by the rule. Also in detail mode, an address column might contain IP addresses in square brackets, for example, [209.165.201.1-209.165.201.30]. These addresses are translated addresses. When an inside host makes a connection to an outside host, the security appliance maps the inside host's address to an address from the pool. After a host creates an outbound connection, the security appliance maintains this address mapping. This address mapping structure is called an xlate, and remains in memory for a period of time.
  - **Service**—Specifies the service and protocol specified by the rule (TCP, UDP, ICMP, or IP).
  - **Action**—Specifies the type of IPSec rule (protect or do not protect).
- **Transform Set**—Displays the transform set for the rule.
- **Peer**—Identifies the IPSec peer.
- **PFS**—Displays Perfect Forward Secrecy settings for the rule.
- **NAT-T Enabled**—Indicates whether NAT Traversal is enabled for the policy.
- **Reverse Route Enabled**—Indicates whether Reverse Route Injection is enabled for the policy.
- **Connection Type**—(Meaningful only for static tunnel policies.) Identifies the connection type for this policy as bidirectional, originate-only, or answer-only).
- **SA Lifetime**—Displays the SA lifetime for the rule.
- **Trustpoint**—Displays the trust point for the policy. This applies to static connections only.
  - **Chain Enabled**—Displays whether the policy transmits the entire trust point chain.
- **IKE Negotiation Mode**—Displays whether IKE negotiations use main or aggressive mode.
- **Description**—(Optional) Specifies a brief description for this rule. For an existing rule, this is the description you typed when you added the rule. An implicit rule includes the following description: “Implicit rule.” To edit the description of any but an implicit rule, right-click this column, and choose Edit Description or double-click the column.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Tunnel Policy (Crypto Map) - Basic

Configuration > VPN > IPSec > IPSec Rules > Add/Edit Rule > Tunnel Policy (Crypto Map) - Basic Tab

Use this pane to define a new Tunnel Policy for an IPSec rule. The values you define here appear in the IPSec Rules table after you click OK. All rules are enabled by default as soon as they appear in the IPSec Rules table.

The Tunnel Policy panel lets you define a tunnel policy that is used to negotiate an IPSec (Phase 2) security association (SA). ASDM captures your configuration edits, but does not save them to the running configuration until you click Apply.

Every tunnel policy must specify a transform set and identify the security appliance interface to which it applies. The transform set identifies the encryption and hash algorithms that perform IPSec encryption and decryption operations. Because not every IPSec peer supports the same algorithms, you might want to specify a number of policies and assign a priority to each. The security appliance then negotiates with the remote IPSec peer to agree on a transform set that both peers support.

Tunnel policies can be *static* or *dynamic*. A static tunnel policy identifies one or more remote IPSec peers or subnetworks to which your security appliance permits IPSec connections. A static policy can be used whether your security appliance initiates the connection or receives a connection request from a remote host. A static policy requires you to enter the information necessary to identify permitted hosts or networks.

A dynamic tunnel policy is used when you cannot or do not want to provide information about remote hosts that are permitted to initiate a connection with the security appliance. If you are only using your security appliance as a VPN client in relation to a remote VPN central-site device, you do not need to configure any dynamic tunnel policies. Dynamic tunnel policies are most useful for allowing remote access clients to initiate a connection to your network through a security appliance acting as the VPN central-site device. A dynamic tunnel policy is useful when the remote access clients have dynamically assigned IP addresses or when you do not want to configure separate policies for a large number of remote access clients.

Fields

- **Interface**—Select the interface name to which this policy applies.
- **Policy Type**—Select the type, static or dynamic, of this tunnel policy.
- **Priority**—Enter the priority of the policy.
- **Transform Set to Be Added**—Select the transform set for the policy and click **Add** to move it to the list of active transform sets. Click **Move Up** or **Move Down** to rearrange the order of the transform sets in the list box. You can add a maximum of 11 transform sets to a crypto map entry or a dynamic crypto map entry.
- **Peer Settings - Optional for Dynamic Crypto Map Entries**—Configure the peer settings for the policy.
  - **Connection Type**—(Meaningful only for static tunnel policies.) Select bidirectional, originate-only, or answer-only to specify the connection type of this policy. For LAN-to-LAN connections, select bidirectional or answer-only (not originate-only). Select answer-only for LAN-to-LAN redundancy.
  - **IP Address of Peer to Be Added**—Enter the IP address of the IPSec peer you are adding.
- **Enable Perfect Forwarding Secrecy**—Check to enable Perfect Forward Secrecy for the policy. PFS is a cryptographic concept where each new key is unrelated to any previous key. In IPSec negotiations, Phase 2 keys are based on Phase 1 keys unless you specify Perfect Forward Secrecy.

- **Diffie-Hellman Group**—When you enable PFS you must also select a Diffie-Hellman group which the security appliance uses to generate session keys. The choices are as follows:
  - Group 1 (768-bits) = Use Perfect Forward Secrecy, and use Diffie-Hellman Group 1 to generate IPSec session keys, where the prime and generator numbers are 768 bits. This option is more secure but requires more processing overhead.
  - Group 2 (1024-bits) = Use Perfect Forward Secrecy, and use Diffie-Hellman Group 2 to generate IPSec session keys, where the prime and generator numbers are 1024 bits. This option is more secure than Group 1 but requires more processing overhead.
  - Group 5 (1536-bits)
  - Group 7 (ECC) = Use Perfect Forward Secrecy, and use Diffie-Hellman Group 7 (ECC) to generate IPSec session keys, where the elliptic curve field size is 163 bits. This option is the fastest and requires the least overhead. It is intended for use with the Movian VPN client, but you can use it with any peers that support Group 7 (ECC).

Tunnel Policy (Crypto Map) - Advanced

Configuration > VPN > IPSec > IPSec Rules > Add/Edit Rule > Tunnel Policy (Crypto Map) - Advanced Tab

Fields

- **Security Association Lifetime** parameters—Configures the duration of a Security Association (SA). This parameter specifies how to measure the lifetime of the IPSec SA keys, which is how long the IPSec SA lasts until it expires and must be renegotiated with new keys.
  - **Time**—Specifies the SA lifetime in terms of hours (hh), minutes (mm) and seconds (ss).
  - **Traffic Volume**—Defines the SA lifetime in terms of kilobytes of traffic. Enter the number of kilobytes of payload data after which the IPSec SA expires. Minimum is 100 KB, default is 10000 KB, maximum is 2147483647 KB.
- **Enable NAT-T**— Enables NAT Traversal (NAT-T) for this policy.
- **Enable Reverse Route Injection**—Enables Reverse Route Injection for this policy.
- **Static Type Only Settings**—Specifies parameters for static tunnel policies.
  - **Trust Point Name**—Selects the trust point to use. If you select something other than None (Use Preshared Keys), which is the default, the Enable entire chain transmission check box becomes active.
  - **Enable entire chain transmission**—Enables transmission of the entire trust point chain.
  - **IKE Negotiation Mode**—Selects the IKE negotiation mode, Main or Aggressive. This parameter sets the mode for exchanging key information and setting up the SAs. It sets the mode that the initiator of the negotiation uses; the responder auto-negotiates. Aggressive Mode is faster, using fewer packets and fewer exchanges, but it does not protect the identity of the communicating parties. Main Mode is slower, using more packets and more exchanges, but it protects the identities of the communicating parties. This mode is more secure and it is the default selection. If you select Aggressive, the Diffie-Hellman Group list becomes active.
  - **Diffie-Hellman Group**—Select the Diffie-Hellman group to apply. The choices are as follows: Group 1 (768-bits), Group 2 (1024-bits), Group 5 (1536-bits), Group 7 (ECC).

Tunnel Policy (Crypto Map) -Traffic Selection

Configuration > VPN > IPSec > IPSec Rules > Add/Edit Rule > Tunnel Policy (Crypto Map) - Traffic Selection Tab

This pane lets you define what traffic to protect.

Fields

- **Interface and Action**
  - **Interface**—Identify the interface for the rule.
  - **Action**—Specify the action for this rule to take. The selections are protect and do not protect.
- **Source** and **Destination**—Specify the IP address, network object group or interface IP address for the source host or network. A rule cannot use the same address as both the source and destination.
  - **IP Address**—Indicates that the parameters that follow specify the interface, IP address, and subnet mask of the source host or network.
  - **Name**—Indicates that the parameters that follow specify the name of the source host or network.
  - **Group**—Indicates that the parameters that follow specify the interface and group name of the source host or network.
  - **Interface**—Selects the interface name for the IP address. This parameter appears when you select the IP Address option button.
  - **IP address**—Specifies the IP address of the interface to which this policy applies. This parameter appears when you select the IP Address option button.
  - **...** —Displays the Select host/network panel, on which you can select an interface and display and select the associated hosts. Your selection appears in the Source Host/Network IP address and Mask boxes on the Add Rule panel. This selection also fills in the Mask field. This parameter appears when you select the IP Address option button.
  - **Mask**—Selects a standard subnet mask to apply to the IP address. This parameter appears when you select the IP Address option button.
  - **Name**—Selects the interface name to use as the source or destination host or network. This parameter appears when you select the Name option button. This is the only parameter associated with this option.
  - **Interface**—Selects the interface name for the IP address. This parameter appears when you select the Group option button.
  - **Group**—Selects the name of the group on the specified interface for the source or destination host or network. If the list contains no entries, you can enter the name of an existing group. This parameter appears when you select the Group option button.
- **Rule Flow Diagram**—Shows the results of your selections for the source and destination host/network group boxes.
- **Protocol and Service**—Specifies protocol and service parameters relevant to this rule.

Note "Any - any" IPSec rules are not allowed. This type of rule would prevent the device and its peer from supporting multiple LAN -to-LAN tunnels.

  - **TCP**—Specifies that this rule applies to TCP connections. This selection also displays the **Source Port** and **Destination Port** group boxes.

- **UDP**—Specifies that this rule applies to UDP connections. This selection also displays the **Source Port** and **Destination Port** group boxes.
- **ICMP**—Specifies that this rule applies to ICMP connections. This selection also displays the **ICMP Type** group box.
- **IP**—Specifies that this rule applies to IP connections. This selection also displays the **IP Protocol** group box.
- **Manage Service Groups**—Displays the Manage Service Groups panel, on which you can add, edit, or delete a group of TCP/UDP services/ports.
- **Source Port** and **Destination Port** —Contains TCP or UDP port parameters, depending on which option button you selected in the Protocol and Service group box.
- **Service**—Indicates that you are specifying parameters for an individual service. Specifies the name of the service and a boolean operator to use when applying the filter.
- **Boolean operator** (unlabeled)—Lists the boolean conditions (equal, not equal, greater than, less than, or range) to use in matching the service specified in the service box.
- **Service** (unlabeled)—Identifies the service (such as https, kerberos, or any) to be matched. If you specified the range service operator this parameter becomes two boxes, into which you enter the start and the end of the range.
- **...** —Displays a list of services from which you can select the service to display in the Service box.
- **Service Group**—Indicates that you are specifying the name of a service group for the source port.
- **Service** (unlabeled)—Selects the service group to use.
- **ICMP Type**—Specifies the ICMP type to use. The default is any. Click the **...** button to display a list of available types.

- **Options**
  - **Time Range**—Specify the name of an existing time range or create a new range.
  - **...** —Displays the Add Time Range pane, on which you can define a new time range.
  - **Please enter the description below (optional)**—Provides space for you to enter a brief description of the rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Pre-Fragmentation

Configuration > VPN > IPSec > Pre-Fragmentation

Use this panel to set the IPSec pre-fragmentation policy and do-not-fragment (DF) bit policy for any interface.

The IPSec pre-fragmentation policy specifies how to treat packets that exceed the maximum transmission unit (MTU) setting when tunneling traffic through the public interface. This feature provides a way to handle cases where a router or NAT device between the security appliance and the client rejects or drops IP fragments. For example, suppose a client wants to FTP get from an FTP server behind a security appliance. The FTP server transmits packets that when encapsulated would exceed the security appliance's MTU size on the public interface. The selected options determine how the security appliance processes these packets. The pre-fragmentation policy applies to all traffic travelling out the security appliance public interface.

The security appliance encapsulates all tunneled packets. After encapsulation, the security appliance fragments packets that exceed the MTU setting before transmitting them through the public interface. This is the default policy. This option works for situations where fragmented packets are allowed through the tunnel without hindrance. For the FTP example, large packets are encapsulated and then fragmented at the IP layer. Intermediate devices may drop fragments or just out-of-order fragments. Load-balancing devices can introduce out-of-order fragments.

When you enable pre-fragmentation, the security appliance fragments tunneled packets that exceed the MTU setting before encapsulating them. If the DF bit on these packets is set, the security appliance clears the DF bit, fragments the packets, and then encapsulates them. This action creates two independent non-fragmented IP packets leaving the public interface and successfully transmits these packets to the peer site by turning the fragments into complete packets to be reassembled at the peer site. In our example, the security appliance overrides the MTU and allows fragmentation by clearing the DF bit.

Note Changing the MTU or the pre-fragmentation option on *any* interface tears down *all* existing connections. For example, if 100 active tunnels terminate on the public interface, and you change the MTU or the pre-fragmentation option on the external interface, all of the active tunnels on the public interface are dropped.

Fields

- **Pre-Fragmentation**—Shows the current pre-fragmentation configuration for every configured interface.
  - **Interface**—Shows the name of each configured interface.
  - **Pre-Fragmentation Enabled**—Shows, for each interface, whether pre-fragmentation is enabled.
  - **DF Bit Policy**—Shows the DF Bit Policy for each interface.
- **Edit**—Displays the Edit IPSec Pre-Fragmentation Policy dialog box.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Edit IPSec Pre-Fragmentation Policy

Configuration > VPN > IPSec > Pre-Fragmentation > Edit IPSec Pre-Fragmentation Policy

Use this panel to modify an existing IPSec pre-fragmentation policy and do-not-fragment (DF) bit policy for an interface selected on the parent panel, **Configuration > VPN > IPSec > Pre-Fragmentation**

Fields

- **Interface**—Identifies the selected interface. You cannot change this parameter using this dialog box.
- **Enable IPSec pre-fragmentation**—Enables or disables IPSec pre-fragmentation. The security appliance fragments tunneled packets that exceed the MTU setting before encapsulating them. If the DF bit on these packets is set, the security appliance clears the DF bit, fragments the packets, and then encapsulates them. This action creates two independent, non-fragmented IP packets leaving the public interface and successfully transmits these packets to the peer site by turning the fragments into complete packets to be reassembled at the peer site.
- **DF Bit Setting Policy**—Selects the do-not-fragment bit policy: Copy, Clear, or Set.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Transform Sets

Configuration > VPN > IPSec > Transform Sets

Use this panel to view and add or edit transform sets. A transform is a set of operations done on a data flow to provide data authentication, data confidentiality, and data compression. For example, one transform is the ESP protocol with 3DES encryption and the HMAC-MD5 authentication algorithm (ESP-3DES-MD5).

Fields

- **Transform Sets**—Shows the configured transform sets.
  - **Name**—Shows the name of the transform sets.
  - **Mode**—Shows the mode, Tunnel, of the transform set. This parameter specifies the mode for applying ESP encryption and authentication; in other words, what part of the original IP packet has ESP applied. Tunnel mode applies ESP encryption and authentication to the entire original IP packet (IP header and data), thus hiding the ultimate source and destination addresses.
  - **ESP Encryption**—Shows the Encapsulating Security Protocol (ESP) encryption algorithms for the transform sets. ESP provides data privacy services, optional data authentication, and anti-replay services. ESP *encapsulates* the data being protected.
  - **ESP Authentication**—Shows the ESP authentication algorithms for the transform sets.
- **Add**—Opens the Add Transform Set dialog box, in which you can add a new transform set.

- **Edit**—Opens the Edit Transform Set dialog box, in which you can modify an existing transform set.
- **Delete**—Removes the selected transform set. There is no confirmation or undo.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Transform Set

Configuration > VPN > IPSec > Transform Sets > Add/Edit Transform Set

Use this panel to add or modify a transform set. A transform is a set of operations done on a data flow to provide data authentication, data confidentiality, and data compression. For example, one transform is the ESP protocol with 3DES encryption and the HMAC-MD5 authentication algorithm (ESP-3DES-MD5).

Fields

- **Set Name**—Specifies a name for this transform set.
- **Properties**—Configures properties for this transform set. These properties appear in the Transform Sets table.
  - **Mode**—Shows the mode, Tunnel, of the transform set. This field shows the mode for applying ESP encryption and authentication; in other words, what part of the original IP packet has ESP applied. Tunnel mode applies ESP encryption and authentication to the entire original IP packet (IP header and data), thus hiding the ultimate source and destination addresses.
  - **ESP Encryption**—Selects the Encapsulating Security Protocol (ESP) encryption algorithms for the transform sets. ESP provides data privacy services, optional data authentication, and anti-replay services. ESP *encapsulates* the data being protected.
  - **ESP Authentication**—Selects the ESP authentication algorithms for the transform sets.

Note The IPSec ESP (Encapsulating Security Payload) protocol provides both encryption and authentication. Packet authentication proves that data comes from whom you think it comes from; it is often referred to as "data integrity."

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Load Balancing

Configuration > VPN > Load Balancing

Note VPN load balancing runs only on security appliance models ASA 5520 and higher.

VPN load balancing requires an active 3DES/AES license. The security appliance checks for the existence of this crypto license before enabling load balancing. If it does not detect an active 3DES or AES license, the security appliance prevents the enabling of load balancing and also prevents internal configuration of 3DES by the load balancing system unless the license permits this usage.

This panel lets you enable load balancing on the security appliance. Enabling load balancing involves:

- Configuring the load-balancing cluster by establishing a common virtual cluster IP address, UDP port (if necessary), and IPSec shared secret for the cluster. These values are identical for every device in the cluster.
- Configuring the participating device by enabling load balancing on the device and defining device-specific properties. These values vary from device to device.

If you have a remote-client configuration in which you are using two or more security appliances connected to the same network to handle remote sessions, you can configure these devices to share their session load. This feature is called *load balancing*. Load balancing directs session traffic to the least loaded device, thus distributing the load among all devices. It makes efficient use of system resources and provides increased performance anodize availability.

Note Load balancing is effective only on remote sessions initiated with the Cisco VPN Client (Release 3.0 and later), the Cisco VPN 3002 Hardware Client (Release 3.5 and later), or the ASA 5505 operating as an Easy VPN Client. All other clients, including LAN-to-LAN connections, can connect to a security appliance on which load balancing is enabled, but the cannot participate in load balancing.

To implement load balancing, you group together logically two or more devices on the same private LAN-to-LAN network into a *virtual cluster.*

All devices in the virtual cluster carry session loads. One device in the virtual cluster, the *virtual cluster master*, directs incoming calls to the other devices, called *secondary devices*. The virtual cluster master monitors all devices in the cluster, keeps track of how busy each is, and distributes the session load accordingly. The role of virtual cluster master is not tied to a physical device; it can shift among devices. For example, if the current virtual cluster master fails, one of the secondary devices in the cluster takes over that role and immediately becomes the new virtual cluster master.

The virtual cluster appears to outside clients as a single *virtual cluster IP address*. This IP address is not tied to a specific physical device. It belongs to the current virtual cluster master; hence, it is virtual. A VPN client attempting to establish a connection connects first to this virtual cluster IP address. The virtual cluster master then sends back to the client the public IP address of the least-loaded available host in the cluster. In a second transaction (transparent to the user) the client connects directly to that host. In this way, the virtual cluster master directs traffic evenly and efficiently across resources.

Note All clients other than the Cisco VPN client, the Cisco VPN 3002 Hardware Client, or the ASA 5505 operating as an Easy VPN Client connect directly to the security appliance as usual; they do not use the virtual cluster IP address.

If a machine in the cluster fails, the terminated sessions can immediately reconnect to the virtual cluster IP address. The virtual cluster master then directs these connections to another active device in the cluster. Should the virtual cluster master itself fail, a secondary device in the cluster immediately and automatically takes over as the new virtual session master. Even if several devices in the cluster fail, users can continue to connect to the cluster as long as any one device in the cluster is up and available.

Prerequisites

Load balancing is disabled by default. You must explicitly enable load balancing.

You must have first configured the public and private interfaces and also have previously configured the the interface to which the virtual cluster IP address refers.

All devices that participate in a cluster must share the same cluster-specific values: IP address, encryption settings, encryption key, and port.

Fields

- **VPN Load Balancing**—Configures virtual cluster device parameters.
  - **Participate in Load Balancing Cluster**—Specifies that this device is a participant in the load-balancing cluster.
  - **VPN Cluster Configuration**—Configures device parameters that must be the same for the entire virtual cluster. All servers in the cluster must have an identical cluster configuration.
  - **Cluster IP Address**—Specifies the single IP address that represents the entire virtual cluster. Choose an IP address that is within the public subnet address range shared by all the security appliances in the virtual cluster.
  - **UDP Port**—Specifies the UDP port for the virtual cluster in which this device is participating. The default value is 9023. If another application is using this port, enter the UDP destination port number you want to use for load balancing.
  - **Enable IPSec Encryption**—Enables or disables IPSec encryption. If you select this check box, you must also specify and verify a shared secret.The security appliances in the virtual cluster communicate via LAN-to-LAN tunnels using IPSec. To ensure that all load-balancing information communicated between the devices is encrypted, select this check box.

Note When using encryption, you must have previously configured the load-balancing inside interface. If that interface is not enabled on the load-balancing inside interface, you get an error message when you try to configure cluster encryption.

If the load-balancing inside interface was enabled when you configured cluster encryption, but was disabled before you configured the participation of the device in the virtual cluster, you get an error message when you select the Participate in Load Balancing Cluster check box, and encryption is not enabled for the cluster.

  - **IPSec Shared Secret**—Specifies the shared secret to between IPSec peers when you have enabled IPSec encryption. The value you enter in the box appears as consecutive asterisk characters.
  - **Verify Secret**—Confirms the shared secret value entered in the IPSec Shared Secret box.
- **VPN Server Configuration**—Configures parameters for this specific device.
  - **Interfaces**—Configures the public and private interfaces and their relevant parameters.
  - **Public**—Specifies the name or IP address of the public interface for this device.

- **Private**—Specifies the name or IP address of the private interface for this device.
- **Priority**—Specifies the priority assigned to this device within the cluster. The range is from 1 to 10. The priority indicates the likelihood of this device becoming the virtual cluster master, either at start-up or when an existing master fails. The higher you set the priority (for example, 10), the more likely this device becomes the virtual cluster master.

Note If the devices in the virtual cluster are powered up at different times, the first device to be powered up assumes the role of virtual cluster master. Because every virtual cluster requires a master, each device in the virtual cluster checks when it is powered-up to ensure that the cluster has a virtual master. If none exists, that device takes on the role. Devices powered up and added to the cluster later become secondary devices. If all the devices in the virtual cluster are powered up simultaneously, the device with the highest priority setting becomes the virtual cluster master. If two or more devices in the virtual cluster are powered up simultaneously, and both have the highest priority setting, the one with the lowest IP address becomes the virtual cluster master.

- **NAT Assigned IP Address**—Specifies the IP address that this device's IP address is translated to by NAT. Enter 0.0.0.0 if NAT is not being used or if the device is not behind a firewall using NAT.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

NAC

Configuration > VPN > NAC

The security appliance uses Extensible Authentication Protocol (EAP) over UDP (EAPoUDP) messaging to validate the posture of remote hosts. Posture validation involves the checking of a remote host for compliancy with safety requirements before the assignment of a network access policy. An Access Control Server must be configured for Network Admission Control before you configure NAC on the security appliance.

The NAC window lets you set attributes that apply to all NAC communications. The following global attributes at the top of the window apply to EAPoUDP messaging between the security appliance and remote hosts:

- Retransmission Timer—The security appliance starts this timer when it sends an EAPoUDP message to the host. A response from the host clears the timer. If the timer expires before the security appliance receives a response, it resends the message. The setting is in seconds. Enter a value in the range 1 to 60. The default setting is 3.
- Hold Timer—The security appliance starts this timer when it places the NAC session for a remote host into a hold state. It places a session in a hold state if it does not receive a response after sending EAPoUDP messages equal to the number of EAPoUDP Retries. The security appliance also starts

this timer after it receives an Access Reject message from the ACS server. When the timer expires, the security appliance tries to initiate a new EAP over UDP association with the remote host. The setting is in seconds. Enter a value in the range 60 to 86400. The default setting is 180.

- EAPoUDP Retries—Number of times the security appliance resends an EAP over UDP message. This attribute limits the number of consecutive retries sent in response to Retransmission Timer expirations. The setting is in seconds. Enter a value in the range 1 to 3. The default setting is 3.
- EAPoUDP Port—Port number for EAP over UDP communication with the Cisco Trust Agent (CTA) on the host. This number must match the port number configured on the CTA. Enter a value in the range 1024 to 65535. The default setting is 21862.

The Clientless Authentication area of the NAC window lets you configure settings for hosts that are not responsive to the EAPoUDP requests. Hosts for which there is no CTA running do not respond to these requests.

- Enable Clientless Authentication—Check to enable clientless authentication. The security appliance sends the configured clientless username and password to the Access Control Server in the form of a user authentication request. The ACS in turn requests the access policy for clientless hosts. If you uncheck this attribute, the security appliance applies the default ACL for clientless hosts.
- Username—Username configured for clientless hosts on the ACS. The default setting is clientless. Enter 1 to 64 ASCII characters, excluding leading and trailing spaces, pound signs (#), question marks (?), single and double quotation marks (“ ” and "), asterisks (*), and angle brackets (< and >).
- Password—Password configured for clientless hosts on the ACS. The default setting is clientless. Enter 4 – 32 ASCII characters.
- Confirm Password—Password configured for clientless hosts on the ACS repeated for validation.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

CHAPTER 28

General

A virtual private network is a network of virtual circuits that carry private traffic over a public network such as the Internet. VPNs can connect two or more LANS, or remote users to a LAN. VPNs provide privacy and security by requiring all users to authenticate and by encrypting all data traffic.

This section describes the general VPN configuration attributes, including the following:

- Client Update
- Default Tunnel Gateway
- Group Policy
- Browse Time Range
- ACL Manager
- Tunnel Group
- VPN System Options
- Zone Labs Integrity Server
- Easy VPN Remote
- Advanced Easy VPN Properties

Client Update

Configuration > VPN > General > Client Update

The **Client Update** window lets administrators at a central location do the following actions:

- Enable the update; specify the types and revision numbers of clients to which the update applies
- Provide a URL or IP address from which to get the update
- In the case of Windows clients, optionally notify users that they should update their VPN client version.

Note The Client Update function at Configuration > VPN > General > Client Update applies only to Windows clients and VPN 3002 hardware clients.

For Windows clients, you can provide a mechanism for users to accomplish that update. For VPN 3002 hardware client users, the update occurs automatically, with no notification. You can apply client updates only to the IPSec remote-access tunnel-group type.

Note If you try to do a client update to an IPSec LAN-to-LAN tunnel group or a WebVPN tunnel group, you do not receive an error message, but no update notification or client update goes to those types of tunnel groups.

To enable client update globally for all clients of a particular client type, use this window. You can also notify all Windows clients that an upgrade is needed and initiate an upgrade on all VPN 3002 hardware clients from this window. To configure the client revisions to which the update applies and the URL or IP address from which to download the update, click Edit.

To configure client update revisions and software update sources for a specific tunnel group, see Configuration > VPN > General > Tunnel Group > Add/Edit > IPSec Tab > Client VPN Software Update Table.

Fields

- Enable Client Update—Enables or disables client update, both globally and for specific tunnel groups. You must enable client update before you can send a client update notification to Windows VPN clients or initiate an automatic update to hardware clients.
- Client Type—Lists the clients to upgrade: software or hardware, and for software clients, all Windows clients or a subset. If you click All Windows Based, do not specify Windows 95, 98 or ME and Windows NT, 2000 or XP individually. The hardware client gets updated with a release of the ASA 5505 software or of the VPN 3002 hardware client.
- VPN Client Revisions—Contains a comma-separated list of software image revisions appropriate for this client. If the user's client revision number matches one of the specified revision numbers, there is no need to update the client, and, for Windows-based clients, the user does not receive an update notification. The following caveats apply:
  - The revision list must include the software version for this update.
  - Your entries must match exactly those on the URL for the VPN client, or the TFTP server for the hardware client.
  - The TFTP server for distributing the hardware client image must be a robust TFTP server.
  - A VPN client user must download an appropriate software version from the listed URL.
  - The VPN 3002 hardware client software is automatically updated via TFTP, with no notification to the user.
- Image URL—Contains the URL or IP address from which to download the software image. This URL must point to a file appropriate for this client. For Windows-based clients, the URL must be in the form: http:// or https://. For hardware clients, the URL must be in the form tftp://.
  - For Windows-based VPN clients: To activate the Launch button on the VPN Client Notification, the URL must include the protocol HTTP or HTTPS and the server address of the site that contains the update. The format of the URL is: http(s)://*server_address:port/directory/filename*. The server address can be either an IP address or a hostname if you have configured a DNS server. For example:

    http://10.10.99.70/vpnclient-win-4.6.Rel-k9.exe

    The directory is optional. You need the port number only if you use ports other than 80 for HTTP or 443 for HTTPS.
  - For the hardware client: The format of the URL is tftp://*server_address/directory/filename*. The server address can be either an IP address or a hostname if you have configured a DNS server. For example:

tftp://10.1.1.1/vpn3002-4.1.Rel-k9.bin

- Edit—Opens the Edit Client Update Entry dialog box, which lets you configure or change client update parameters. See Edit Client Update Entry.
- Live Client Update—Sends an upgrade notification message to all currently connected VPN clients or selected tunnel group(s).
  - Tunnel Group—Selects all or specific tunnel group(s) for updating.
  - Update Now—Immediately sends an upgrade notification containing a URL specifying where to retrieve the updated software to the currently connected Windows VPN clients in the selected tunnel group or all connected tunnel groups. The message includes the location from which to download the new version of software. The administrator for that VPN client can then retrieve the new software version and update the VPN client software.

    For VPN 3002 hardware clients, the upgrade proceeds automatically, with no notification.

    You must check Enable Client Update in the window for the upgrade to work. Clients that are not connected receive the upgrade notification or automatically upgrade the next time they log on.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Edit Client Update Entry

Configuration > VPN > General > Client Update > Edit Client Update Entry

The Edit Client Update dialog box lets you change information about VPN client revisions and URLs for the indicated client types. The clients must be running one of the revisions specified for the indicated client type. If not, the clients are notified that an upgrade is required.

Fields

- Client Type—(*Display-only*) Displays the client type selected for editing.
- VPN Client Revisions—Lets you type a comma-separated list of software or firmware images appropriate for this client. If the user's client revision number matches one of the specified revision numbers, there is no need to update the client. If the client is not running a software version on the list, an update is in order. The user of a Windows-based VPN client must download an appropriate software version from the listed URL. The VPN 3002 hardware client software is automatically updated via TFTP.
- Image URL—Lets you type the URL for the software/firmware image. This URL must point to a file appropriate for this client.
  - For a Windows-based VPN client, the URL must include the protocol HTTP or HTTPS and the server address of the site that contains the update. The format of the URL is: http(s)://*server_address:port/directory/filename*. The server address can be either an IP address or a hostname if you have configured a DNS server. For example:

```
http://10.10.99.70/vpnclient-win-4.6.Rel-k9.exe
```

The directory is optional. You need the port number only if you use ports other than 80 for HTTP or 443 for HTTPS.

– For the hardware client: The format of the URL is tftp://*server_address/directory/filename*. The server address can be either an IP address or a hostname if you have configured a DNS server. For example:

tftp://10.1.1.1/vpn3002-4.1.Rel-k9.bin

The directory is optional.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Default Tunnel Gateway

Configuration > VPN > General > Default Tunnel Gateway

To configure the default tunnel gateway, click the Static Route link in this window. The Configuration > Routing > Routing > Static Route window opens.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Group Policy

Configuration > VPN > General > Group Policy

The Group Policy window lets you manage VPN group policies. A VPN group policy is a collection of user-oriented attribute/value pairs stored either internally on the device or externally on a RADIUS or LDAP server. Configuring the VPN group policy lets users inherit attributes that you have not configured at the individual group or username level. By default, VPN users have no group policy association. The group policy information is used by VPN tunnel groups and user accounts.

The “child” windows, tabs, and dialog boxes let you configure the default group parameters. These parameters are those that are most likely to be common across all groups and users, and they streamline the configuration task. Groups can “inherit” parameters from this default group, and users can “inherit” parameters from their group or the default group. You can override these parameters as you configure groups and users.

If you click the Add dialog box, a small menu appears giving you the option to create a new internal group policy, or an external group policy that is stored externally on a RADIUS or LDAP server. Both the Add Internal Group Policy window and the Edit Group Policy window include six tabbed sections. If you click the WebVPN tab, you expose six additional tabs. Click each tab to display its parameters. As you move from tab to tab, the security appliance retains your settings. When you have finished setting parameters on all tabbed sections, click OK or Cancel.

In these dialog boxes, you configure the following kinds of parameters:

- General Parameters: Protocols, filtering, connection settings, and servers.
- IPSec Parameters: IP Security tunneling protocol parameters and client access rules.
- Client Configuration Parameters: Banner, password storage, split-tunneling policy, default domain name, IPSec over UDP, backup servers.
- Client FW Parameters: VPN Client personal firewall requirements.
- Hardware Client Parameters: Interactive hardware client and individual user authentication; network extension mode.
- WebVPN Parameters: SSL VPN access.

Before configuring these parameters, you should configure:

- Access hours.
- Rules and filters.
- IPSec Security Associations.
- Network lists for filtering and split tunneling
- User authentication servers, and specifically the internal authentication server.

Fields

- Group Policy—Contains a table listing the currently configured group policies and Add, Edit, and Delete buttons to help you manage VPN group policies.
  - Name—Lists the name of the currently configured group policies.
  - Type—Lists the type of each currently configured group policy.
  - Tunneling Protocol—Lists the tunneling protocol that each currently configured group policy uses.
  - AAA Server Group—Lists the AAA server group, if any, to which each currently configured group policy pertains.
  - Add—Displays the Add Group Policy dialog box, which lets you add a new AAA group policy to the list. This screen includes seven tabbed sections. Click each tab to display its parameters. As you move from tab to tab, ASDM retains your settings. When you have finished setting parameters on all tabbed sections, click Apply or Cancel.
  - Edit—Displays the Edit Group Policy dialog box, which lets you modify an existing AAA group policy.
  - Delete—Lets you remove a AAA group policy from the list. There is no confirmation or undo.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit External Group Policy

Configuration > VPN > General > Group Policy > Add > External Group Policy > Add or Edit External Group Policy

The Add or Edit External Group Policy dialog box lets you configure an external group policy.

Fields

- Name—Identifies the group policy to be added or changed. For Edit External Group Policy, this field is display-only.
- Server Group—Lists the available server groups to which to apply this policy.
- Password—Specifies the password for this server group policy.
- New—Opens a dialog box that lets you select whether to create a new RADIUS server group or a new LDAP server group. Either of these options opens the Add AAA Server Group dialog box.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add AAA Server Group

Configuration > VPN > General > Group Policy > Add/Edit > External Group Policy > New > RADIUS Server Group/New LDAP Server Group > Add AAA Server Group

The Add AAA Server Group dialog box lets you configure a new AAA server group. The Accounting Mode attribute applies only to RADIUS and TACACS+ protocols.

Fields

- Server Group—Specifies the name of the server group.
- Protocol—(*Display only*) Indicates whether this is a RADIUS or an LDAP server group.
- Accounting Mode—Indicates whether to use simultaneous or single accounting mode. In single mode, the security appliance sends accounting data to only one server. In simultaneous mode, the security appliance sends accounting data to all servers in the group. The Accounting Mode attribute applies only to RADIUS and TACACS+ protocols.

- Reactivation Mode—Specifies the method by which failed servers are reactivated: Depletion or Timed reactivation mode. In Depletion mode, failed servers are reactivated only after all of the servers in the group become inactive. In Timed mode, failed servers are reactivated after 30 seconds of down time.
- Dead Time—Specifies, for depletion mode, the number of minutes (0 through 1440) that must elapse between the disabling of the last server in the group and the subsequent re-enabling of all servers. The default value is 10 minutes. This field is not available for timed mode.
- Max Failed Attempts— Specifies the number (an integer in the range 1 through 5) of failed connection attempts allowed before declaring a nonresponsive server inactive. The default value is 3 attempts.

Add/Edit Internal Group Policy > General Tab

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > General Tab

The Add or Edit Group Policy window, General tab, lets you specify tunneling protocols, filters, connection settings, and servers for the group policy being added or modified. For each of the fields on this window, checking the Inherit check box lets the corresponding setting take its value from the default group policy.

Fields

- Inherit—(Multiple instances) Indicates that the corresponding setting takes its value from the default group policy. This is the default value for all of the attributes on this tab.
- Tunneling Protocols—Specifies the tunneling protocols that this group can use. Users can use only the selected protocols. The choices are as follows:
  - IPSec—IP Security Protocol. Regarded as the most secure protocol, IPSec provides the most complete architecture for VPN tunnels. Both LAN-to-LAN (peer-to-peer) connections and client-to-LAN connections can use IPSec.
  - WebVPN—VPN via SSL/TLS. Uses a web browser to establish a secure remote-access tunnel to a security appliance; requires neither a software nor hardware client. WebVPN can provide easy access to a broad range of enterprise resources, including corporate websites, web-enabled applications, NT/AD file share (web-enabled), e-mail, and other TCP-based applications from almost any computer that can reach HTTPS Internet sites.
  - L2TP over IPSec—Allows remote users with VPN clients provided with several common PC and mobile PC operating systems to establish secure connections over the public IP network to the security appliance and private corporate networks. L2TP uses PPP over UDP (port 1701) to tunnel the data. The security appliance must be configured for IPSec transport mode.

  Note If no protocol is selected, an error message appears.

- Filter—Specifies the filter to use, or whether to inherit the value from the group policy. Filters consist of rules that determine whether to allow or reject tunneled data packets coming through the security appliance, based on criteria such as source address, destination address, and protocol. To configure filters and rules, see the Configuration > Features > VPN > VPN General > Group Policy window.
- Manage—Displays the ACL Manager window, with which you can add, edit, and delete Access Control Lists (ACLs) and Extended Access Control Lists (ACEs). For more information about the ACL Manager, see the online Help for that window.

- Connection Settings—Specifies the connection settings parameters.
  - Access Hours—If the Inherit check box is not selected, you can select the name of an existing access hours policy, if any, applied to this user or create a new access hours policy. The default value is Inherit, or, if the Inherit check box is not selected, the default value is --Unrestricted--.
  - Manage—Opens the Browse Time Range dialog box, on which you can add, edit, or delete a time range.
  - Simultaneous Logins—If the Inherit check box is not selected, this parameter specifies the maximum number of simultaneous logins allowed for this user. The default value is 3. The minimum value is 0, which disables login and prevents user access.

Note While there is no maximum limit, allowing several simultaneous connections might compromise security and affect performance.

  - Maximum Connect Time—If the Inherit check box is not selected, this parameter specifies the maximum user connection time in minutes. At the end of this time, the system terminates the connection. The minimum is 1 minute, and the maximum is 35791394 minutes (over 4000 years). To allow unlimited connection time, select Unlimited (the default).
  - Idle Timeout—If the Inherit check box is not selected, this parameter specifies this user's idle timeout period in minutes. If there is no communication activity on the user's connection in this period, the system terminates the connection. The minimum time is 1 minute, and the maximum time is 10080 minutes. The default is 30 minutes. To allow unlimited connection time, select Unlimited. This value does not apply to WebVPN users.
- Servers—Configures DNS and WINS servers, and DHCP Scope.
  - DNS Servers—Specifies the DNS servers to use. If you deselect Inherit, you can specify the primary and secondary DNS servers in their respective boxes.
  - WINS Servers—Specifies the WINS servers to use. If you deselect Inherit, you can specify the primary and secondary WINS servers in their respective boxes.
  - DHCP Scope—Specifies the DHCP scope; that is, the range of IP addresses the security appliance DHCP server should use to assign addresses to users of this group policy. If you deselect Inherit, you can enter the scope in the box.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Browse Time Range

You can get to this panel through various paths.

Use the Browse Time Range dialog box to add, edit, or delete a time range. A time range is a reusable component that defines starting and ending times that can be applied to a group policy. After defining a time range, you can select the time range and apply it to different options that require scheduling. For example, you can attach an access list to a time range to restrict access to the security appliance. A time range consists of a start time, an end time, and optional recurring (that is, periodic) entries. For more information about time ranges, see the online Help for the Add or Edit Time Range dialog box.

Fields

- Add—Opens the Add Time Range dialog box, on which you can create a new time range.

Note Creating a time range does not restrict access to the device.

- Edit—Opens the Edit Time Range dialog box, on which you can modify an existing time range. This button is active only when you have selected an existing time range from the Browse Time Range table.
- Delete—Removes a selected time range from the Browse Time Range table. There is no confirmation or undo of this action.
- Name—Specifies the name of the time range.
- Start Time—Specifies when the time range begins.
- End Time—Specifies when the time range ends.
- Recurring Entries—Specifies further constraints of active time of the range within the start and stop time specified.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Time Range

You can get to this panel through various paths.

The Add or Edit Time Range dialog box lets you configure a new time range.

Fields

- Time Range Name—Specifies the name that you want to assign to this time range.
- Start Time—Defines the time when you want the time range to start.
  - Start now—Specifies that the time range starts immediately.
  - Start at—Selects the month, day, year, hour, and minute at which you want the time range to start.
- End Time—Defines the time when you want the time range to end.

  - Never end—Specifies that the time range has no defined end point.
  - End at (inclusive)—Selects the month, day, year, hour, and minute at which you want the time range to end.
- Recurring Time Ranges—Constrains the active time of this time range within the start and end times when the time range is active. For example, if the start time is start now and the end time is never end, and you want the time range to be effective every weekday, Monday through Friday, from 8:00 AM to 5:00 PM, you could configure a recurring time range, specifying that it is to be active weekdays from 08:00 through 17:00, inclusive.
- Add—Opens the Add Recurring Time Range dialog box, on which you can configure a recurring time range.
- Edit—Opens the Edit Recurring Time Range dialog box, on which you can modify a selected recurring time range.
- Delete—Removes a selected recurring time range.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Add/Edit Recurring Time Range

You can get to this panel through various paths.

The Add or Edit Recurring Time Range dialog box lets you configure or modify a recurring time range.

Fields

- Specify days of the week and times on which this recurring range will be active—Makes available the options in the Days of the week area. For example, use this option when you want the time range to be active only every Monday through Thursday, from 08:00 through 16:59.
  - Days of the week—Select the days that you want to include in this recurring time range. Possible options are: Every day, Weekdays, Weekends, and On these days of the week. For the last of these, you can select a check box for each day that you want included in the range.
  - Daily Start Time—Specifies the hour and minute, in 24-hour format, when you want the recurring time range to be active on each selected day.
  - Daily End Time (inclusive)—Specifies the hour and minute, in 24-hour format, when you want the recurring time range to end on each selected day.
- Specify a weekly interval when this recurring range will be active—Makes available the options in the Weekly Interval area. The range extends inclusively through the end time. All times in this area are in 24-hour format. For example, use this option when you want the time range to be active continuously from Monday at 8:00 AM through Friday at 4:30 PM.
  - From—Selects the day, hour, and minute when you want the weekly time range to start.
  - Through—Selects the day, hour, and minute when you want the weekly time range to end.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

ACL Manager

You can get to this panel through various paths.

The ACL Manager dialog box lets you define access control lists (ACLs) to control the access of a specific host or network to another host/network, including the protocol or port that can be used.

You can configure ACLs (Access Control Lists) to apply to user sessions. These are filters that permit or deny user access to specific networks, subnets, hosts, and web servers.

- If you do not define any filters, all connections are permitted.
- The security appliance supports only an inbound ACL on an interface.
- At the end of each ACL, there is an implicit, unwritten rule that denies all traffic that is not permitted. If traffic is not explicitly permitted by an access control entry (ACE), the security appliance denies it. ACEs are referred to as rules in this topic.

Standard ACL Tab

This pane provides summary information about standard ACLs, and lets you add or edit ACLs and ACEs.

Fields

- Add—Lets you add a new ACL. When you highlight an existing ACL, it lets you add a new ACE for that ACL.
- Edit—Opens the Edit ACE dialog box, on which you can change an existing access control list rule.
- Delete—Removes an ACL or ACE. There is no confirmation or undo.
- Move Up/Move Down—Changes the position of a rule in the ACL Manager table.
- Cut—Removes the selection from the ACL Manager table and places it on the clipboard.
- Copy—Places a copy of the selection on the clipboard.
- Paste—Opens the Paste ACE dialog box, on which you can create a new ACL rule from an existing rule.
- No—Indicates the order of evaluation for the rule. Implicit rules are not numbered, but are represented by a hyphen.
- Address—Displays the IP address or URL of the application or service to which the ACE applies.
- Action—Specifies whether this filter permits or denies traffic flow.
- Description—Shows the description you typed when you added the rule. An implicit rule includes the following description: “Implicit outbound rule.”

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Extended ACL Tab

This pane provides summary information about extended ACLs, and lets you add or edit ACLs and ACEs.

Fields

- Add—Lets you add a new ACL. When you highlight an existing ACL, it lets you add a new ACE for that ACL.
- Edit—Opens the Edit ACE dialog box, on which you can change an existing access control list rule.
- Delete—Removes an ACL or ACE. There is no confirmation or undo.
- Move Up/Move Down—Changes the position of a rule in the ACL Manager table.
- Cut—Removes the selection from the ACL Manager table and places it on the clipboard.
- Copy—Places a copy of the selection on the clipboard.
- Paste—Opens the Paste ACE dialog box, on which you can create a new ACL rule from an existing rule.
- No—Indicates the order of evaluation for the rule. Implicit rules are not numbered, but are represented by a hyphen.
- Enabled—Enables or disables a rule. Implicit rules cannot be disabled.
- Source—Specifies the IP addresses (Host/Network) that are permitted or denied to send traffic to the IP addresses listed in the Destination column. In detail mode (see the Show Detail radio button), an address column might contain an interface name with the word any, such as inside: any. This means that any host on the inside interface is affected by the rule.
- Destination—Specifies the IP addresses (Host/Network) that are permitted or denied to send traffic to the IP addresses listed in the Source column. An address column might contain an interface name with the word any, such as outside: any. This means that any host on the outside interface is affected by the rule. An address column might also contain IP addresses; for example 209.165.201.1-209.165.201.30. These addresses are translated addresses. When an inside host makes a connection to an outside host, the firewall maps the address of the inside host to an address from the pool. After a host creates an outbound connection, the firewall maintains this address mapping. The address mapping structure is called an xlate, and remains in memory for a period of time. During this time, outside hosts can initiate connections to the inside host using the translated address from the pool, if allowed by the ACL. Normally, outside-to-inside connections require a static translation so that the inside host always uses the same IP address.
- Service—Names the service and protocol specified by the rule.
- Action—Specifies whether this filter permits or denies traffic flow.

- Logging —Shows the logging level and the interval in seconds between log messages (if you enable logging for the ACL). To set logging options, including enabling and disabling logging, right-click this column, and choose Edit Log Option. The Log Options window appears.
- Time—Specifies the name of the time range to be applied in this rule.
- Description—Shows the description you typed when you added the rule. An implicit rule includes the following description: “Implicit outbound rule.”

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit/Paste ACE

ACL Manager > Add/Edit/Paste Extended ACE

The Add/Edit/Paste ACE dialog box lets you create a new extended access list rule, or modify an existing rule. The Paste option becomes available only when you cut or copy a rule.

Fields

- Action—Determines the action type of the new rule. Select either permit or deny.
  - Permit—Permits all matching traffic.
  - Deny—Denies all matching traffic.
- Source/Destination—Specifies the source or destination type and, depending on that type, the other relevant parameters describing the source or destination host/network IP Address. Possible values are: any, IP address, Network Object Group, and Interface IP. The availability of subsequent fields depends upon the value of the Type field:
  - any—Specifies that the source or destination host/network can be any type. For this value of the Type field, there are no additional fields in the Source or Destination area.
  - IP Address—Specifies the source or destination host or network IP address. With this selection, the IP Address, ellipsis button, and Netmask fields become available. Select an IP address or host name from the drop-down list in the IP Address field or click the ellipsis (...) button to browse for an IP address or name. Select a network mask from the drop-down list.
  - Network Object Group—Specifies the name of the network object group. Select a name from the drop-down list or click the ellipsis (...) button to browse for a network object group name.
  - Interface IP—Specifies the interface on which the host or network resides. Select an interface from the drop-down list. The default values are inside and outside. There is no browse function.
- Protocol and Service—Specifies the protocol and service to which this ACE filter applies. Service groups let you identify multiple non-contiguous port numbers that you want the ACL to match. For example, if you want to filter HTTP, FTP, and port numbers 5, 8, and 9, define a service group that includes all these ports. Without service groups, you would have to create a separate rule for each port.

You can create service groups for TCP, UDP, TCP-UDP, ICMP, and other protocols. A service group with the TCP-UDP protocol contains services, ports, and ranges that might use either the TCP or UDP protocol.

  - Protocol—Selects the protocol to which this rule applies. Possible values are ip, tcp, udp, icmp, and other. The remaining available fields in the Protocol and Service area depend upon the protocol you select. The next few bullets describe the consequences of each of these selections:
  - Protocol: TCP and UDP—Selects the TCP/UDP protocol for the rule. The Source Port and Destination Port areas allow you to specify the ports that the ACL uses to match packets.
  - Source Port/Destination Port—(*Available only for TCP and UDP protocols*) Specifies an operator and a port number, a range of ports, or a well-known service name from a list of services, such as HTTP or FTP. The operator list specifies how the ACL matches the port. Choose one of the following operators: = (equals the port number), not = (does not equal the port number), > (greater than the port number), < (less than the port number), range (equal to one of the port numbers in the range).
  - Group—(*Available only for TCP and UDP protocols*) Selects a source port service group. The Browse (...) button opens the Browse Source Port or Browse Destination Port dialog box.
  - Protocol: ICMP—Lets you select an ICMP type or ICMP group from a preconfigured list or browse (...) for an ICMP group. The Browse button opens the Browse ICMP dialog box.
  - Protocol: IP—Specifies the IP protocol for the rule in the IP protocol box. No other fields are available when you make this selection.
  - Protocol: Other—Lets you select a protocol from a drop-down list, select a protocol group from a drop-down list, or browse for a protocol group. The Browse (...) button opens the Browse Other dialog box.

- Rule Flow Diagram—(*Display only*) Provides a graphical representation of the configured rule flow. This same diagram appears on the ACL Manager dialog box unless you explicitly close that display.
- Options—Sets optional features for this rule, including logging parameters, time ranges, and description.
  - Logging—Enables or disables logging or specifies the use of the default logging settings. If logging is enabled, the Syslog Level and Log Interval fields become available.
  - Syslog Level—Selects the level of logging activity. The default is Informational.
  - Log Interval—Specifies the interval for permit and deny logging. The default is 300 seconds. The range is 1 through 6000 seconds.
  - Time Range—Selects the name of the time range to use with this rule. The default is (any). Click the Browse (...) button to open the Browse Time Range dialog box to select or add a time range.
  - Description—(*Optional*) Provides a brief description of this rule. A description line can be up to 100 characters long, but you can break a description into multiple lines.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Browse Source/Destination Address

ACL Manager > Add/Edit Extended Access List Rule > Source or Destination > Browse button

The Browse Source or Destination Address dialog box lets you select an object to use a s a source or destination for this rule.

Fields

- Type—Determines the type of object to use as the source or destination for this rule. Selections are IP Address Objects, IP Names, Network Object Groups, and All. The contents of the table following this field change, depending upon your selection.
- Source/Destination Object Table—Displays the objects from which you can select a source or destination object. If you select All in the type field, each category of object appears under its own heading. The table has the following headings:
  - Name—Displays the network name (which may be an IP address) for each object.
  - IP address—Displays the IP address of each object.
  - Netmask—Displays the network mask to use with each object.
  - Description—Displays the description entered in the Add/Edit/Paste Extended Access List Rule dialog box.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Browse Source/Destination Port

ACL Manager > Add/Edit Extended Access List Rule > Protocol and Service > Protocol: tcp or udp >Source or Destination Port > Group option > Browse button

The Browse Source or Destination Port dialog box lets you select a source or destination port for this protocol in this rule.

Fields

- Add—Opens the Add TCP Service Group dialog box, on which you can configure a new TCP service group.
- Find—Opens the Filter field.
- Filter/Clear—Specifies a filter criterion that you can use to search for items in the Name list, thus displaying only those items that match that criterion. When you make an entry in the Filter field, the Filter button becomes active. Clicking the Filter button performs the search. After you perform the search, the Filter button is dimmed, and the Clear button becomes active. Clicking the Clear button clears the filter field and dims the Clear button.

- Type—Determines the type of object to use as the source or destination for this rule. Selections are IP Address Objects, IP Names, Network Object Groups, and All. The contents of the table following this field change, depending upon your selection.
- Name—Lists the predefined protocols and service groups for your selection.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add TCP Service Group

ACL Manager > Add/Edit Extended Access List Rule > Protocol and Service > Protocol: tcp or udp >Source or Destination Port > Group option > Browse button > Browse Source or Destination Port > Add button

The Add TCP Service Group dialog box lets you configure a new a TCP service group or port to add to the browsable source or destination port list for this protocol in this rule. Selecting a member of either the Members not in Group or the Members in Group list activates the Add and Remove buttons.

Fields

- Group Name—Specifies the name of the new TCP service group.
- Description—(Optional) Provides a brief description of this group.
- Members not in Group—Presents the option to select either a service/service group or a port number to add to the Members in Group list.
- Service/Service Group—Selects the option to select the name of a TCP service or service group to add to the Members in Group list.
- Port #—Selects the option to specify a range of port numbers to add to the Members in Group list.
- Add—Moves a selected item from the Members not in Group list to the Members in Group list.
- Remove—Moves a selected item from the Members in Group list to the Members not in Group list.
- Members in Group—Lists the members already configured in this service group.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Browse ICMP

ACL Manager > Add/Edit Extended Access List Rule > Protocol and Service > Protocol: icmp >ICMP > Group option > Browse button

The Browse ICMP dialog box lets you select an ICMP group for this rule.

Fields

- Add—Opens the Add ICMP Group dialog box, on which you can configure a new TCP service group.
- Find—Opens the Filter field.
- Filter/Clear—Specifies a filter criterion that you can use to search for items in the Name list, thus displaying only those items that match that criterion. When you make an entry in the Filter field, the Filter button becomes active. Clicking the Filter button performs the search. After you perform the search, the Filter button is dimmed, and the Clear button becomes active. Clicking the Clear button clears the filter field and dims the Clear button.
- Type—Determines the type of object to use as the ICMP group for this rule. Selections are IP Address Objects, IP Names, Network Object Groups, and All. The contents of the table following this field change, depending upon your selection.
- Name—Lists the predefined ICMP groups for your selection.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add ICMP Group

ACL Manager > Add/Edit Extended Access List Rule > Protocol and Service > Protocol: icmp >ICMP > Group option > Browse button > Browse ICMP > Add button

The Add ICMP Group dialog box lets you configure a new a ICMP group by name or by number to add to the browsable ICMP list for this protocol in this rule. Selecting a member of either the Members not in Group or the Members in Group list activates the Add and Remove buttons.

Fields

- Group Name—Specifies the name of the new TCP service group.
- Description—(Optional) Provides a brief description of this group.
- Members not in Group—Presents the option to select either an ICMP type/ICMP group or an ICMP number to add to the Members in Group list.
- ICMP Type/ICMP Group—Selects the option to select the name of an ICMP group to add to the Members in Group list.
- ICMP #—Selects the option to specify an ICMP member by number to add to the Members in Group list.

- Add—Moves a selected item from the Members not in Group list to the Members in Group list.
- Remove—Moves a selected item from the Members in Group list to the Members not in Group list.
- Members in Group—Lists the members already configured in this service group.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Browse Other

ACL Manager > Add/Edit Extended Access List Rule > Protocol and Service > Protocol: other >Other > Group option > Browse button

The Browse Other dialog box lets you select a protocol group for this rule.

Fields

- Add—Opens the Add Protocol Group dialog box, on which you can configure a new service group.
- Find—Opens the Filter field.
- Filter/Clear—Specifies a filter criterion that you can use to search for items in the Name list, thus displaying only those items that match that criterion. When you make an entry in the Filter field, the Filter button becomes active. Clicking the Filter button performs the search. After you perform the search, the Filter button is dimmed, and the Clear button becomes active. Clicking the Clear button clears the filter field and dims the Clear button.
- Type—Determines the type of object to use as the protocol group for this rule. Selections are IP Address Objects, IP Names, Network Object Groups, and All. The contents of the table following this field change, depending upon your selection.
- Name—Lists the predefined protocol groups for your selection.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add Protocol Group

ACL Manager > Add/Edit Extended Access List Rule > Protocol and Service > Protocol: other > Group option > Browse button > Browse Other > Add button

The Add Protocol Group dialog box lets you configure a new a protocol group by name or by number to add to the browsable protocol list for this rule. Selecting a member of either the Members not in Group or the Members in Group list activates the Add and Remove buttons.

Fields

- Group Name—Specifies the name of the new TCP service group.
- Description—(Optional) Provides a brief description of this group.
- Members not in Group—Presents the option to select either a protocol/protocol group or a protocol number to add to the Members in Group list.
- Protocol/Protocol Group—Selects the option to select the name of a protocol or protocol group to add to the Members in Group list.
- Protocol #—Selects the option to specify a protocol by number to add to the Members in Group list.
- Add—Moves a selected item from the Members not in Group list to the Members in Group list.
- Remove—Moves a selected item from the Members in Group list to the Members not in Group list.
- Members in Group—Lists the members already configured in this service group.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Internal Group Policy > IPSec Tab

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > IPSec Tab

The Add or Edit Group Policy window, IPSec tab, lets you specify tunneling protocols, filters, connection settings, and servers for the group policy being added or modified.

Fields

- Inherit—(Multiple instances) Indicates that the corresponding setting takes its value from the default group policy. This is the default option for all attributes on this tab.
- Re-Authentication on IKE Re-key—Enables or disables reauthentication when IKE re-key occurs, unless the Inherit check box is selected.
- IP Compression—Enables or disables IP Compression, unless the Inherit check box is selected.
- Perfect Forward Secrecy—Enables or disables perfect forward secrecy (PFS), unless the Inherit check box is selected. PFS ensures that the key for a given IPSec SA was not derived from any other secret (like some other keys). In other words, if someone were to break a key, PFS ensures that the attacker would not be able to derive any other key. If PFS were not enabled, someone could hypothetically break the IKE SA secret key, copy all the IPSec protected data, and then use knowledge of the IKE SA secret to compromise the IPSec SAs set up by this IKE SA. With PFS, breaking IKE would not give an attacker immediate access to IPSec. The attacker would have to break each IPSec SA individually.

- Tunnel Group Lock—Enables locking the tunnel group you select from the list, unless the Inherit check box or the value None is selected.
- Client Access Rules—Lets you configure up to 25 client access rules. If you deselect the Inherit check box, the Add, Edit, and Delete buttons become active and the following column headings appear in the table:
  - Priority—Shows the priority for this rule.
  - Action—Specifies whether this rule permits or denies access.
  - Client Type—Specifies the type of VPN client to which this rule applies, software or hardware, and for software clients, all Windows clients or a subset.
  - VPN Client Version—Specifies the version or versions of the VPN client to which this rule applies. This box contains a comma-separated list of software or firmware images appropriate for this client.
- Add—Adds a new rule for an IPSec group policy. This button is active only if the Inherit check box is deselected.
- Edit—Modifies an existing rule for an IPSec group policy. This button is active only if the Inherit check box is deselected.
- Delete—Removes an existing rule for an IPSec group policy. This button is active only if the **Inherit** check box is deselected. There is no confirmation or undo.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Client Access Rule

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > IPSec > Add or Edit Client Access Rule

Spawned from the Add or Edit Group Policy windows, IPSec tab, the Add or Edit Client Access Rule window adds a new client access rule for an IPSec group policy or modifies an existing rule.

Fields

- Priority—Shows the priority for this rule.
- Action—Specifies whether this rule permits or denies access.
- VPN Client Type—Specifies the type of VPN client to which this rule applies, software or hardware, and for software clients, all Windows clients or a subset. Some common values for VPN Client Type include VPN 3002, PIX, Linux, * (matches all client types), Win9x (matches Windows 95, Windows 98, and Windows ME), and WinNT (matches Windows NT, Windows 2000, and Windows XP). If you choose *, do not configure individual Windows types such as Windows NT.
- VPN Client Version—Specifies the version or versions of the VPN client to which this rule applies. This box contains a comma-separated list of software or firmware images appropriate for this client. The following caveats apply:

- You must specify the software version for this client. You can specify * to match any version.
- Your entries must match exactly those on the URL for the VPN client, or the TFTP server for the VPN 3002.
- The TFTP server for distributing the hardware client image must be a robust TFTP server.
- If the client is already running a software version on the list, it does not need a software update. If the client is not running a software version on the list, an update is in order.
- A VPN client user must download an appropriate software version from the listed URL.
- The VPN 3002 hardware client software is automatically updated via TFTP.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Internal Group Policy > Client Configuration Tab

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Client Configuration Tab

The Add or Edit Group Policy window, Client Configuration tab contains three tabs that let you configure general client parameters, Cisco client parameters, and Microsoft client parameters.

For information about the individual tabs, see the following links:

- Add/Edit Internal Group Policy > Client Configuration Tab > General Client Parameters Tab
- Add/Edit Internal Group Policy > Client Configuration Tab > Cisco Client Parameters Tab
- Add/Edit Internal Group Policy > Client Configuration Tab > Microsoft Client Parameters Tab

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Internal Group Policy > Client Configuration Tab > General Client Parameters Tab

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Client Configuration Tab >General Client Parameters Tab

This tab configures client attributes that are common across both Cisco and Microsoft clients, including the banner text, default domain, split tunnel parameters, and address pools.

Fields

- Inherit—(Multiple instances) Indicates that the corresponding setting takes its value from the default group policy. Deselecting the Inherit check box makes other options available for the parameter. This is the default option for all attributes on this tab.
- Banner—Specifies whether to inherit the banner from the default group policy or enter new banner text. For more information, see View/Config Banner
- Edit Banner—Displays the View/Config Banner dialog box, in which you can enter banner text, up to 500 characters.
- Default Domain—Specifies whether to inherit the default domain from the default group policy or use a new default domain specified in the field.
- Split Tunnel DNS Names (space delimited)—Specifies whether to inherit the split-tunnel DNS names or from the default group policy or specify a new name or list of names in the field.
- Split Tunnel Policy—Specifies whether to inherit the split-tunnel policy from the default group policy or select a policy from the menu. The menu options are to tunnel all networks, tunnel those in the network list below, or exclude those in the network list below.
- Split Tunnel Network List—Specifies whether to inherit the split-tunnel network list from the default group policy or select from the drop-down list.
- Manage—Opens the ACL Manager dialog box, on which you can manage standard and extended access control lists.
- Address Pools—Configures the address pools available through this group policy.
  - Available Pools—Specifies a list of address pools for allocating addresses to remote clients. Deselecting the Inherit check box with no address pools in the Assigned Pools list indicates that no address pools are configured and disables inheritance from other sources of group policy.
  - Add—Moves the name of an address pool from the Available Pools list to the Assigned Pools list.
  - Remove—Moves the name of an address pool from the Assigned Pools list to the Available Pools list.
  - Assigned Pools (up to 6 entries)—Lists the address pools you have added to the assigned pools list. The address-pools settings in this table override the local pool settings in the group. You can specify a list of up to six local address pools to use for local address allocation. The order in which you specify the pools is significant. The security appliance allocates addresses from these pools in the order in which the pools appear in this command.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

View/Config Banner

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Client Configuration > Edit Banner > View/Config Banner

The View/Config Banner dialog box lets you enter into the text box up to 500 characters of text to be displayed as a banner for the specified client.

Note A carriage return/line feed, created by pressing Enter, counts as 2 characters.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Internal Group Policy > Client Configuration Tab > Cisco Client Parameters Tab

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Client Configuration Tab >Cisco Client Parameters Tab

This tab configures client attributes that are specific to Cisco clients, including password storage, enabling or disabling IPSec over UDP and setting the UDP port number, and configuring IPSec backup servers.

Fields

- Store Password on Client System—Enables or disables storing the password on the client system.

Note Storing the password on a client system can constitute a potential security risk.

- IPSec over UDP—Enables or disables using IPSec over UDP.
- IPSec over UDP Port—Specifies the UDP port to use for IPSec over UDP.
- IPSec Backup Servers—Activates the Server Configuration and Server IP Addresses fields, so you can specify the UDP backup servers to use if these values are not inherited.
- Server Configuration—Lists the server configuration options to use as an IPSec backup server. The available options are: Keep Client Configuration (the default), Use the Backup Servers Below, and Clear Client Configuration.
- Server Addresses (space delimited)—Specifies the IP addresses of the IPSec backup servers. This field is available only when the value of the Server Configuration selection is Use the Backup Servers Below.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Internal Group Policy > Client Configuration Tab > Microsoft Client Parameters Tab

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Client Configuration Tab >Microsoft Client Parameters Tab

This tab configures client attributes that are specific to Microsoft clients, specifically, proxy server parameters for Microsoft Internet Explorer.

Fields

- Proxy Server Policy—Configures the Microsoft Internet Explorer browser proxy actions ("methods") for a client PC.
  - Do not modify client proxy settings—Leaves the HTTP browser proxy server setting in Internet Explorer unchanged for this client PC.
  - Do not use proxy—Disables the HTTP proxy setting in Internet Explorer for the client PC.
  - Auto-detect proxy—Enables the use of automatic proxy server detection in Internet Explorer for the client PC.
  - Use proxy server settings specified below—Sets the HTTP proxy server setting in Internet Explorer to use the value configured in the Proxy Server Name or IP Address field.
- Proxy Server Settings—Configures the proxy server parameters for Microsoft clients using Microsoft Internet Explorer.
  - Proxy Server Name or IP Address—Specifies the IP address or name of an Microsoft Internet Explorer server that is applied for this client PC.

Note ASDM lets you configure the proxy server name or IP address. To configure the optional port to use, as well as the server, you must use the **msie-proxy server** command in group-policy configuration mode.

  - Bypass Proxy Server for Local Addresses— Configures Microsoft Internet Explorer browser proxy local-bypass settings for a client PC. Select Yes to enable local bypass or No to disable local bypass.
  - Proxy Server Exception List—Configures Microsoft Internet Explorer browser proxy exception list settings for a local bypass on the client PC. Enter the list of addresses that you do not want to have accessed through a proxy server. This list corresponds to the Exceptions box in the Proxy Settings dialog box in Internet Explorer.
  - Name or IP Address (use * as a wildcard)—Specifies the IP address or name of an MSIE server that is applied for this client PC.
  - Add—Add the specified name or IP address to the Proxy Server Exceptions list.
  - Delete—Remove the specified name or IP address from the Proxy server Exceptions list.

  - Proxy Server Exceptions—Lists the server names and IP addresses that you want to exclude from proxy server access. This list corresponds to the Exceptions box in the Proxy Settings dialog box in Internet Explorer.

- DHCP Intercept—Enables or disables DHCP Intercept. DHCP Intercept lets Microsoft XP clients use split-tunneling with the security appliance. The security appliance replies directly to the Microsoft Windows XP client DHCP Inform message, providing that client with the subnet mask, domain name, and classless static routes for the tunnel IP address. For Windows clients prior to XP, DHCP Intercept provides the domain name and subnet mask. This is useful in environments in which using a DHCP server is not advantageous.

Note A Microsoft XP anomaly results in the corruption of domain names if split tunnel options exceed 255 bytes. To avoid this problem, the security appliance limits the number of routes it sends to 27 to 40 routes, with the number of routes dependent on the classes of the routes.

  - Intercept DHCP Configure Message—Specifies whether to inherit the DHCP intercept policy from the group policy or to enable (Yes) or disable (No) DHCP policy.
  - Subnet Mask (optional)—Selects the subnet mask from the drop-down list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Standard Access List Rule

ACL Manager > Add or Edit Standard Access List Rule

The Add/Edit Standard Access List Rule dialog box lets you create a new rule, or modify an existing rule.

Fields

- Action—Determines the action type of the new rule. Select either permit or deny.
  - Permit—Permits all matching traffic.
  - Deny—Denies all matching traffic.
- Host/Network IP Address—Identifies the networks by IP address.
  - IP address—The IP address of the host or network.
  - Mask—The subnet mask of the host or network
- Description—(Optional) Enter a description of the access rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Internal Group Policy > Client Firewall Tab

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Client Firewall Tab

The Add or Edit Group Policy window, Client Firewall tab, lets you configure firewall settings for VPN clients for the group policy being added or modified.

Note Only VPN clients running Microsoft Windows can use these firewall features. They are currently not available to hardware clients or other (non-Windows) software clients.

A *firewall* isolates and protects a computer from the Internet by inspecting each inbound and outbound individual packet of data to determine whether to allow or drop it. Firewalls provide extra security if remote users in a group have split tunneling configured. In this case, the firewall protects the user's PC, and thereby the corporate network, from intrusions by way of the Internet or the user's local LAN. Remote users connecting to the security appliance with the VPN client can choose the appropriate firewall option.

In the first scenario, a remote user has a personal firewall installed on the PC. The VPN client enforces firewall policy defined on the local firewall, and it monitors that firewall to make sure it is running. If the firewall stops running, the VPN client drops the connection to the security appliance. (This firewall enforcement mechanism is called *Are You There (AYT)*, because the VPN client monitors the firewall by sending it periodic "are you there?" messages; if no reply comes, the VPN client knows the firewall is down and terminates its connection to the security appliance.) The network administrator might configure these PC firewalls originally, but with this approach, each user can customize his or her own configuration.

In the second scenario, you might prefer to enforce a centralized firewall policy for personal firewalls on VPN client PCs. A common example would be to block Internet traffic to remote PCs in a group using split tunneling. This approach protects the PCs, and therefore the central site, from intrusions from the Internet while tunnels are established. This firewall scenario is called *push policy* or *Central Protection Policy (CPP)*. On the security appliance, you create a set of traffic management rules to enforce on the VPN client, associate those rules with a filter, and designate that filter as the firewall policy. The security appliance pushes this policy down to the VPN client. The VPN client then in turn passes the policy to the local firewall, which enforces it.

Fields

- Inherit—Determines whether the group policy obtains its client firewall setting from the default group policy. This option is the default setting. When set, it overrides the remaining attributes in this tab and dims their names.
- Client Firewall Attributes—Specifies the client firewall attributes, including what type of firewall (if any) is implemented and the firewall policy for that firewall.

- Firewall Setting—Lists whether a firewall exists, and if so, whether it is required or optional. If you select No Firewall (the default), none of the remaining fields on this window are active. If you want users in this group to be firewall-protected, select either the Firewall Required or Firewall Optional setting.

  If you select Firewall Required, all users in this group must use the designated firewall. The security appliance drops any session that attempts to connect without the designated, supported firewall installed and running. In this case, the security appliance notifies the VPN client that its firewall configuration does not match.

Note If you require a firewall for a group, make sure the group does not include any clients other than Windows VPN clients. Any other clients in the group (including ASA 5505 in client mode and VPN 3002 hardware clients) are unable to connect.

  If you have remote users in this group who do not yet have firewall capacity, choose Firewall Optional. The Firewall Optional setting allows all the users in the group to connect. Those who have a firewall can use it; users that connect without a firewall receive a warning message. This setting is useful if you are creating a group in which some users have firewall support and others do not—for example, you may have a group that is in gradual transition, in which some members have set up firewall capacity and others have not yet done so.

- Firewall Type—Lists firewalls from several vendors, including Cisco. If you select Custom Firewall, the fields under Custom Firewall become active. The firewall you designate must correlate with the firewall policies available. The specific firewall you configure determines which firewall policy options are supported.
- Custom Firewall—Specifies the vendor ID, Product ID and description for the custom firewall.
  - Vendor ID—Specifies the vendor of the custom firewall for this group policy.
  - Product ID—Specifies the product or model name of the custom firewall being configured for this group policy.
  - Description—(Optional) Describes the custom firewall.
- Firewall Policy—Specifies the type and source for the custom firewall policy.
  - Policy defined by remote firewall (AYT)—Specifies that the firewall policy is defined by the remote firewall (Are You There). Policy defined by remote firewall (AYT) means that remote users in this group have firewalls located on their PCs. The local firewall enforces the firewall policy on the VPN client. The security appliance allows VPN clients in this group to connect only if they have the designated firewall installed and running. If the designated firewall is not running, the connection fails. Once the connection is established, the VPN client polls the firewall every 30 seconds to make sure that it is still running. If the firewall stops running, the VPN client ends the session.
  - Policy pushed (CPP)—Specifies that the policy is pushed from the peer. If you select this option, the Inbound Traffic Policy and Outbound Traffic Policy lists and the Manage button become active.The security appliance enforces on the VPN clients in this group the traffic management rules defined by the filter you choose from the Policy Pushed (CPP) drop-down menu. The choices available on the menu are filters defined on this security appliance, including the default filters. Keep in mind that the security appliance pushes these rules down to the VPN client, so you should create and define these rules relative to the VPN client, not the security appliance. For example, “in” and “out” refer to traffic coming into the VPN client or going outbound from the VPN client. If the VPN client also has a local firewall, the policy pushed from the security appliance works with the policy of the local firewall. Any packet that is blocked by the rules of either firewall is dropped.

- Inbound Traffic Policy—Lists the available push policies for inbound traffic.
- Outbound Traffic Policy—Lists the available push policies for outbound traffic.
- Manage—Displays the ACL Manager window, on which you can configure Access Control Lists (ACLs).

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Internal Group Policy > Hardware Client Tab

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Hardware Client Tab

The Add or Edit Group Policy window, Hardware Client tab, lets you configure settings for the VPN 3002 hardware client for the group policy being added or modified. The Hardware Client tab parameters do not pertain to the ASA 5505 in client mode.

Fields

- Inherit—(Multiple instances) Indicates that the corresponding setting takes its value from the default group policy, rather than from the explicit specifications that follow. This is the default setting for all attributes in this tab.
- Require Interactive Client Authentication—Enables or disables the requirement for interactive client authentication. This parameter is disabled by default. Interactive hardware client authentication provides additional security by requiring the VPN 3002 to authenticate with a username and password that you enter manually each time the VPN 3002 initiates a tunnel. With this feature enabled, the VPN 3002 does not have a saved username and password. When you enter the username and password, the VPN 3002 sends these credentials to the security appliance to which it connects. The security appliance facilitates authentication, on either the internal or an external authentication server. If the username and password are valid, the tunnel is established.

  When you enable interactive hardware client authentication for a group, the security appliance pushes that policy to the VPN 3002s in the group. If you have previously set a username and password on the VPN 3002, the software deletes them from the configuration file. When you try to connect, the software prompts you for a username and password.

  If, on the security appliance, you subsequently disable interactive hardware authentication for the group, it is enabled locally on the VPN 3002s, and the software continues to prompt for a username and password. This lets the VPN 3002 connect, even though it lacks a saved username and password, and the security appliance has disabled interactive hardware client authentication. If you subsequently configure a username and password, the feature is disabled, and the prompt no longer appears. The VPN 3002 connects to the security appliance using the saved username and password.
- Require Individual User Authentication—Enables or disables the requirement for individual user authentication for users behind ASA 5505 in client mode or the VPN 3002 hardware client in the group. To display a banner to hardware clients in a group, individual user authentication must be enabled. This parameter is disabled by default.

Individual user authentication protects the central site from access by unauthorized persons on the private network of the hardware client. When you enable individual user authentication, each user that connects through a hardware client must open a web browser and manually enter a valid username and password to access the network behind the security appliance, even though the tunnel already exists.

Note You cannot use the command-line interface to log in if user authentication is enabled. You must use a browser.

If you have a default home page on the remote network behind the security appliance, or if you direct the browser to a website on the remote network behind the security appliance, the hardware client directs the browser to the proper pages for user login. When you successfully log in, the browser displays the page you originally entered.

If you try to access resources on the network behind the security appliance that are not web-based, for example, e-mail, the connection fails until you authenticate using a browser.

To authenticate, you must enter the IP address for the private interface of the hardware client in the browser Location or Address field. The browser then displays the login screen for the hardware client. To authenticate, click the Connect/Login Status button.

One user can log in for a maximum of four sessions simultaneously. Individual users authenticate according to the order of authentication servers configured for a group.

- User Authentication Idle Timeout—Configures a user timeout period. The security appliance terminates the connection if it does not receive user traffic during this period. You can specify that the timeout period is a specific number of minutes or unlimited.
  - Unlimited—Specifies that the connection never times out. This option prevents inheriting a value from a default or specified group policy.
  - Minutes—Specifies the timeout period in minutes. Use an integer between 1 and 35791394. The default value is Unlimited.
- Cisco IP Phone Bypass—Lets Cisco IP phones bypass the interactive individual user authentication processes. If enabled, interactive hardware client authentication remains in effect. Cisco IP Phone Bypass is disabled by default.

Note You must configure the ASA 5505 in client mode or the VPN 3002 hardware client to use network extension mode for IP phone connections.

- LEAP Bypass—Lets LEAP packets from Cisco wireless devices bypass the individual user authentication processes (if enabled). LEAP Bypass lets LEAP packets from devices behind a hardware client travel across a VPN tunnel *prior* to individual user authentication. This lets workstations using Cisco wireless access point devices establish LEAP authentication. Then they authenticate again per individual user authentication (if enabled). LEAP Bypass is disabled by default.

Note This feature does not work as intended if you enable interactive hardware client authentication.

IEEE 802.1X is a standard for authentication on wired and wireless networks. It provides wireless LANs with strong mutual authentication between clients and authentication servers, which can provide dynamic per-user, per-session wireless encryption privacy (WEP) keys, removing administrative burdens and security issues that are present with static WEP keys.

Cisco Systems has developed an 802.1X wireless authentication type called Cisco LEAP. LEAP implements mutual authentication between a wireless client on one side of a connection and a RADIUS server on the other side. The credentials used for authentication, including a password, are always encrypted before they are transmitted over the wireless medium.

Note Cisco LEAP authenticates wireless clients to RADIUS servers. It does not include RADIUS accounting services.

LEAP users behind a hardware client have a circular dilemma: they cannot negotiate LEAP authentication because they cannot send their credentials to the RADIUS server behind the central site device over the tunnel. The reason they cannot send their credentials over the tunnel is that they have not authenticated on the wireless network. To solve this problem, LEAP Bypass lets LEAP packets, and only LEAP packets, traverse the tunnel to authenticate the wireless connection to a RADIUS server before individual users authenticate. Then the users proceed with individual user authentication.

LEAP Bypass works as intended under the following conditions:

- The interactive unit authentication feature (intended for wired devices) must be disabled. If interactive unit authentication is enabled, a non-LEAP (wired) device must authenticate the hardware client before LEAP devices can connect using that tunnel.
- Individual user authentication is enabled (if it is not, you do not need LEAP Bypass).
- Access points in the wireless environment must be Cisco Aironet Access Points. The wireless NIC cards for PCs can be other brands.
- The Cisco Aironet Access Point must be running Cisco Discovery Protocol (CDP).
- The ASA 5505 or VPN 3002 can operate in either client mode or network extension mode.
- LEAP packets travel over the tunnel to a RADIUS server via ports 1645 or 1812.

Note Allowing any unauthenticated traffic to traverse the tunnel might pose a security risk.

- Allow Network Extension Mode—Restricts the use of network extension mode on the hardware client. Select the option to let hardware clients use network extension mode. Network extension mode is required for the hardware client to support IP phone connections, because the Call Manager can communicate only with actual IP addresses.

Note If you disable network extension mode, the default setting, the hardware client can connect to this security appliance in PAT mode only. If you disallow network extension mode here, be careful to configure all hardware clients in a group for PAT mode. If a hardware client is configured to use network extension mode and the security appliance to which it connects disables network extension mode, the hardware client attempts to connect every 4 seconds, and every attempt is rejected. In this situation, the hardware client puts an unnecessary processing load on the security appliance to which it connects; large numbers of hardware clients that are misconfigured in this way reduces the ability of the security appliance to provide service.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Internal Group Policy > NAC Tab

Configuration > VPN > General > Group Policy > Add/Edit Internal Group Policy > NAC Tab

The Add or Edit Internal Group Policy window, NAC tab, lets you configure Network Admission Control settings for the default group policy or an alternative group policy.

Fields

- Inherit—(Multiple instances) Indicates that the corresponding setting takes its value from the default group policy, rather than from the explicit specifications that follow. This is the default setting for all attributes in this tab.
- Enable NAC—Requires posture validation for remote access. If the remote computer passes the validation checks, the ACS server downloads the access policy for the security appliance to enforce. The default setting is Disable.
- Status Query Timer—The security appliance starts this timer after each successful posture validation and status query response. The expiration of this timer triggers a query for changes in the host posture, referred to as a *status query*. Enter the number of seconds in the range 30 to 1800. The default setting is 300.
- Revalidation Timer—The security appliance starts this timer after each successful posture validation. The expiration of this timer triggers the next unconditional posture validation. The security appliance maintains posture validation during revalidation. The default group policy becomes effective if the Access Control Server is unavailable during posture validation or revalidation. Enter the interval in seconds between each successful posture validation. The range is 300 to 86400. The default setting is 36000.
- Default ACL— (Optional) The security appliance applies the security policy associated with the selected ACL if posture validation fails. Select None or select an extended ACL in the list. The default setting is None. If the setting is None and posture validation fails, the security appliance applies the default group policy.

  Use the Manage button to populate the drop-down list and view the configuration of the ACLs in the list.
- Manage— Opens the ACL Manager dialog box. Click to view, enable, disable, and delete standard ACLs and the ACEs in each ACL. The list next to the Default ACL attribute displays the ACLs.
- Posture Validation Exception List—Displays one or more attributes that exempt remote computers from posture validation. At minimum, each entry lists the operating system and an Enabled setting of Yes or No. An optional filter identifies an ACL used to match additional attributes of the remote computer. An entry that consists of an operating system and a filter requires the remote computer to match both to be exempt from posture validation. The security appliance ignores the entry if the Enabled setting is set to No.
- Add—Adds an entry to the Posture Validation Exception list.
- Edit—Modifies an entry in the Posture Validation Exception list.
- Delete—Removes an entry from the Posture Validation Exception list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Posture Validation Exception

Configuration > VPN > General > Group Policy > Add/Edit Internal Group Policy > NAC tab > Add/Edit

The Add/Edit Posture Validation Exception dialog window lets you exempt remote computers from posture validation, based on their operating system and other optional attributes that match a filter.

- Operating System—Choose the operating system of the remote computer. If the computer is running this operating system, it is exempt from posture validation. The default setting is blank.
- Enable—The security appliance checks the remote computer for the attribute settings displayed in this window only if you check Enabled. Otherwise, it ignores the attribute settings. The default setting is unchecked.
- Filter— (Optional) Use to apply an ACL to filter the traffic if the operating system of the computer matches the value of the Operating System attribute.
- Manage— Opens the ACL Manager dialog box. Click to view, enable, disable, and delete standard ACLs and the ACEs in each ACL. The list next to the Default ACL attribute displays the ACLs. Use this button to populate the list next to the Filter attribute.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

WebVPN Tab > Functions Tab

You can get to this panel through various paths.

The WebVPN tab > Functions tab lets you configure the features available to WebVPN users. The interface visible to these WebVPN users varies depending on the values you set here. Users see a customized home page that includes only those features that you enable.

- Inherit —Indicates that the corresponding setting takes its value from the default group policy, rather than from the explicit specifications that follow.
- Enable URL entry—Places the URL entry box on the home page. If this feature is enabled, users can enter web addresses in the URL entry box, and use WebVPN to access those websites.

Using WebVPN does not ensure that communication with every site is secure. WebVPN ensures the security of data transmission between the remote user's PC or workstation and the security appliance on the corporate network. If a user then accesses a non-HTTPS web resource (located on the Internet or on the internal network), the communication from the corporate security appliance to the destination web server is not secured.

In a WebVPN connection, the security appliance acts as a proxy between the end user's web browser and target web servers. When a WebVPN user connects to an SSL-enabled web server, the security appliance establishes a secure connection and validates the server's SSL certificate. The end user's browser never receives the presented certificate, so therefore cannot examine and validate the certificate. The current implementation of WebVPN does not permit communication with sites that present expired certificates. Neither does the security appliance perform trusted CA certificate validation. Therefore, WebVPN users cannot analyze the certificate an SSL-enabled web-server presents before communicating with it.

To limit Internet access for WebVPN users, deselect the Enable URL Entry field. This prevents WebVPN users from surfing the Web during a WebVPN connection.

- Enable file server access—Enables Windows file access (SMB/CIFS files only) through HTTPS. When this box is checked, users can access Windows files on the network. If you enable only this parameter for WebVPN file sharing, users can access only servers that you configure in the Servers and URLs area. To let users access servers directly or to browse servers on the network, see the Enable file server entry and Enable file server browsing attribute descriptions.

  With this box checked, users can download, edit, delete, rename, and move files. They can also add files and folders.

  Shares must also be configured for user access on the applicable Windows servers. Users might have to be authenticated before accessing files, depending on network requirements.

  File access, server/domain access, and browsing require that you configure a WINS server or a master browser, typically on the same network as the security appliance, or reachable from that network. The WINS server or master browser provides the security appliance with an list of the resources on the network. You cannot use a DNS server instead.

Note File access is not supported in an Active Native Directory environment when used with Dynamic DNS. It is supported if used with a WINS server.

- Enable file server entry—Places the file server entry box on the portal page. File server access must be enabled.

  With this box checked, users can enter pathnames to Windows files directly. They can download, edit, delete, rename, and move files. They can also add files and folders. Again, shares must also be configured for user access on the applicable Windows servers. Users might have to be authenticated before accessing files, depending on network requirements.

- Enable file server browsing—Lets users browse the Windows network for domains/workgroups, servers and shares. File server access must be enabled.

  With this box checked, users can select domains and workgroups, and can browse servers and shares within those domains. Shares must also be configured for user access on the applicable Windows servers. Users may need to be authenticated before accessing servers, according to network requirements.

- Enable auto applet download—Lets users automatically download and start the port forwarding java applet upon WebVPN login. Disabled by default, you can enable this feature only if port forwarding, Outlook/Exchange proxy, or HTTP proxy is also enabled. You can also enable auto applet download in the default group policy (DfltGrpPolicy) or in user-defined group policies.

- Enable port forwarding—WebVPN Port Forwarding provides access for remote users in the group to client/server applications that communicate over known, fixed TCP/IP ports. Remote users can use client applications that are installed on their local PC and securely access a remote server that supports that application. Cisco has tested the following applications: Windows Terminal Services, Telnet, Secure FTP (FTP over SSH), Perforce, Outlook Express, and Lotus Notes. Other TCP-based applications may also work, but Cisco has not tested them.

Note Port Forwarding does not work with some SSL/TLS versions.

With this box checked users can access client/server applications by mapping TCP ports on the local and remote systems.

Caution Make sure Sun Microsystems Java™ Runtime Environment (JRE) 1.5.x is installed on the remote computers to support port forwarding (application access) and digital certificates. If JRE 1.4.x is running and the user authenticates with a digital certificate, the application fails to start because JRE cannot access the web browser's certificate store.

- Enable Outlook/Exchange proxy—Enables the use of the Outlook/Exchange e-mail proxy.
- Apply Web-type ACL—Applies the WebVPN access control list defined for the users of this group.
- Enable HTTP Proxy—Enables the forwarding of an HTTP applet proxy to the client. The proxy is useful for technologies that interfere with proper content transformation, such as Java, ActiveX, and Flash. It bypasses mangling while ensuring the continued use of the security appliance.
  The forwarded proxy modifies the browser's old proxy configuration automatically and redirects all HTTP and HTTPS requests to the new proxy configuration. It supports virtually all client side technologies, including HTML, CSS, JavaScript, VBScript, ActiveX, and Java. The only browser it supports is Microsoft Internet Explorer.
- Enable Citrix MetaFrame—Enables support for terminal services from a MetaFrame Application Server to the client. This attribute lets the security appliance act as a secure gateway within a secure Citrix configuration. These services provide users with access to MetaFrame applications through a standard Web browser.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Group Policy > WebVPN Tab > Content Filtering Tab

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Web VPN Tab > Content Filtering Tab

- The Add or Edit Group Policy window, WebVPN tab, Content Filtering tab, lets you configure the security appliance to block or remove the parts of websites that use Java or Active X, scripts, display images, and deliver cookies. By default, these parameters are disabled, which means that no filtering occurs.

Fields

- Inherit—Determines whether this group policy inherits its content filtering values from the default group policy. This option is the default setting. When this attribute is checked, the remaining attributes are dim, indicating that you cannot set them.
- Filter Java/ActiveX—Removes <applet>, <embed> and <object> tags from HTML.
- Filter scripts—Removes <script> tags from HTML.
- Filter images—Removes <img> tags from HTML. Removing images dramatically speeds the delivery of web pages.
- Filter cookies from images—Removes cookies that are delivered with images. This may preserve user privacy, because advertisers use cookies to track visitors.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Group Policy > WebVPN Tab > Homepage Tab

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Web VPN Tab > Homepage Tab

The Add or Edit Group Policy window, WebVPN tab, Homepage tab, lets you configure what, if any, home page to use and specify any customizations (such as color, logo, and so on) that you want to apply to it. It does not define the home page customization.

Fields

- Inherit—Indicates that the corresponding setting takes its value from the default group policy, rather than from the explicit specifications that follow. This is the default setting for both the Webpage Customization and Custom Homepage attributes.
- Webpage Customization—Specifies whether to inherit the webpage customizations from the default group policy, to apply an existing customization (selected from a list), or to create a new customization.
- New—Opens the Add Customization Object dialog box, on which you can create and configure a new customization to apply to the GUI pages that the user sees.
- Custom Homepage—Specifies whether to inherit the home page from the default group policy, use an existing URL as the home page, or use no home page.
- Specify URL—Indicates that the subsequent fields specify the protocol, either http or https, and the URL of the Web page to use as the home page, as follows:

  - Protocol—Indicates whether to use http or https as the connection protocol for the home page.
  - :// field—Specifies the URL of the Web page to use as the home page.
- Use none—Specifies that no home page is configured.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Group Policy > WebVPN Tab > Port Forwarding Tab

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Web VPN Tab > Port Forwarding Tab

The Add or Edit Group Policy window, WebVPN tab, Port Forwarding tab, lets you configure port forwarding parameters.

Fields

- Inherit—(Multiple instances) If checked, this option specifies that the default group policy sets the value of the associated attribute. This option is the default setting for both the Port Forwarding List and Applet Name attributes.
- Port Forwarding List—Specifies whether to inherit the port forwarding list from the default group policy, select one from the list, or create a new port forwarding list.
- New—Opens the Add Port Forwarding List window, on which you can add a new port forwarding list. See the description of the Add Port Forwarding List window.
- Applet Name—Specifies whether to inherit the applet name or to use the name specified in the field. Specify this name to identify port forwarding to end users. The name you configure appears in the end user interface as a hotlink. When users click this link, a Java applet opens a window that displays a table that lists and provides access to port forwarding applications that you configure for these users. The default applet name is Application Access.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Port Forwarding List

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Web VPN Tab > Port Forwarding Tab > New button > Add or Edit Port Forwarding List

The Add Port Forwarding List dialog box lets you specify the name of a port forwarding list and displays a list of configured port forwarding entries.

Fields

- List Name—Assigns a name to the port forwarding list you want to add.
- Local TCP Port—Lists the local TCP port for each entry in the port forwarding list.
- Remote Server—Lists the remote server for each entry in the port forwarding list.
- Remote TCP Port—Lists the remote TCP port for each entry in the port forwarding list.
- Description—(*Optional*) Lists a description, up to 64 characters long, for each entry in the port forwarding list.
- Add—Opens the Add Port Forwarding Entry dialog box, on which you can configure a new port forwarding entry.
- Edit—Opens the Edit Port Forwarding Entry dialog box, on which you can modify an existing port forwarding entry.
- Delete—Removes a selected port forwarding entry from the port forwarding list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Port Forwarding Entry

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Web VPN Tab > Port Forwarding Tab > New button > Add Port Forwarding List > Add or Edit

The Add or Edit Port Forwarding Entry dialog box lets you specify the name of a port forwarding list and displays a list of configured port forwarding entries.

Fields

- Local TCP Port—Specifies the local TCP port for this port forwarding list entry.
- Remote Server—Specifies the remote server for this port forwarding list entry.
- Remote TCP Port—Specifies the remote TCP port for this port forwarding list entry.
- Description—(*Optional*) Specifies a description, up to 64 characters, for this port forwarding list entry.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Group Policy > WebVPN Tab > Other Tab

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Web VPN Tab > Other Tab

The Add or Edit Group Policy window, WebVPN tab, Other tab, lets you configure servers and URL lists and the Web-type ACL ID.

Fields

- Inherit—(Multiple instances) Indicates that the corresponding setting takes its value from the default group policy, rather than from the explicit specifications that follow.
- Servers and URL Lists—Specifies whether to inherit the list of Servers and URLs, to select an existing list, or to create a new list.
- New—Displays a dialog box in which you can add a new server or URL to the list.
- Web-Type ACL ID—Specifies whether to inherit the web-type ACL ID, select the identifier of an existing Web-Type ACL to use, or add or modify a web-type ACL.
- Manage—Opens the ACL Manager dialog box on which you can manage web-type ACLs.
- SSO Server—Specifies whether to inherit the single-sign-on server setting, to select an existing SSO server from the list, or to add a new SSO server.
- New—Opens the Add SSO Server dialog box, on which you can configure a new server for the list.
- HTTP Compression—Specifies whether to inherit the HTTP Compression setting from the default group, or explicitly to enable or disable HTTP compression.
- Keepalive Ignore—Specifies whether to inherit the maximum transaction size from the default group or sets the upper limit of the HTTP/HTTPS traffic, per transaction, to ignore. The range is 0 through 900 KB.
- Deny Message—Lets you inherit, specify, or remove the message to be sent to remote users who log in to WebVPN successfully, but have no VPN privileges, as follows:
  - Check Inherit to inherit from the default group the message to be sent to remote users who log in to WebVPN successfully, but have no VPN privileges.
  - Uncheck and erase the text in the field, to *not* send a message to remote users who log into WebVPN successfully, but have no VPN privileges.
  - Uncheck, and create or modify the message (up to 490 characters long) in the field, to be sent to remote users who log in to WebVPN successfully, but have no VPN privileges. The default message is as follows: “Login was successful, but because certain criteria have not been met or due to some specific group policy, you do not have permission to use any of the VPN features. Contact your IT administrator for more information.”

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Server and URL List

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Web VPN Tab > Other Tab > Add or Edit Server and URL List

The Add or Edit Server and URL List dialog box lets you add, edit, delete, and order the items in the designated URL list.

Fields

- List Name—Specifies the name of the list to be added or selects the name of the list to be modified or deleted.
- URL Display Name—Specifies the URL name displayed to the user.
- URL—Specifies the actual URL associated with the display name.
- Add—Opens the Add Server or URL dialog box, on which you can configure a new server or URL and display name.
- Edit—Opens the Edit Server or URL dialog box, on which you can configure a new server or URL and display name.
- Delete—Removes the selected item from the server and URL list. There is no confirmation or undo.
- Move Up/Move Down—Changes the position of the selected item in the server and URL list.

Add/Edit Server or URL

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Web VPN Tab > Other Tab > Add or Edit Server and URL

The Add or Edit Server or URL dialog box lets you add or edit, delete, and order the items in the designated URL list.

Fields

- URL Display Name—Specifies the URL name displayed to the user.
- URL—Specifies the actual URL associated with the display name.

Add/Edit Group Policy > WebVPN Tab > SSL VPN Client Tab

Configuration > VPN > General > Group Policy > Add/Edit > Internal Group Policy > Web VPN Tab > SSL VPN Client Tab

The Add or Edit Group Policy window, WebVPN tab, SSL VPN Client tab, lets you configure the security appliance to download SSL VPN clients (SVCs) to remote computer.

SVC is a VPN tunneling technology that gives remote users the benefits of an IPSec VPN client without the need for network administrators to install and configure IPSec VPN clients on remote computers. The SVC uses the SSL encryption that is already present on the remote computer as well as the WebVPN login and authentication of the security appliance.

To establish an SVC session, the remote user enters the IP address of a WebVPN interface of the security appliance in the browser, and the browser connects to that interface and displays the WebVPN login screen. If the user satisfies the login and authentication, and the security appliance identifies the user as *requiring* the SVC, the security appliance downloads the SVC to the remote computer. If the security appliance identifies the user as having the *option* to use the SVC, the security appliance downloads the SVC to the remote computer while presenting a link on the user screen to skip the SVC installation.

After downloading, the SVC installs and configures itself, and then the SVC either remains or uninstalls itself (depending on the configuration) from the remote computer when the connection terminates.

The security appliance might have several unique SVC images residing in cache memory for different remote computer operating systems. When the user attempts to connect, the security appliance can consecutively download portions of these images to the remote computer until the image and operating system match, at which point it downloads the entire SVC. You can order the SVC images to minimize connection setup time, with the first image downloaded representing the most commonly-encountered remote computer operating system.

Fields

- Inherit—(Multiple instances) Indicates that the corresponding setting takes its value from the default group policy, rather than from the explicit specifications that follow. This is the default setting for all attributes in this tab.
- Use SSL VPN Client—Specifies whether to inherit the value of this attribute from the default group policy, or when to use the SSL VPN Client: always, optionally, or never.
- Keep Installer on Client System—Enables (Yes) permanent SVC installation or disables (No) the automatic uninstalling feature of the SVC. The SVC remains installed on the remote computer for subsequent SVC connections, reducing the SVC connection time for the remote user.
- Compression—Enables or disables compression on the SVC connection.

  SVC compression increases the communications performance between the security appliance and the SVC by reducing the size of the packets being transferred.
- Keepalive Messages—Adjusts the frequency of keepalive messages, in the range of 15 to 600 seconds.

  You can adjust the frequency of keepalive messages to ensure that an SVC connection through a proxy, firewall, or NAT device remains open, even if the device limits the time that the connection can be idle. Adjusting the frequency also ensures that the SVC does not disconnect and reconnect when the remote user is not actively running a socket-based application, such as Microsoft Outlook or Microsoft Internet Explorer.
- Key Renegotiation Settings—When the security appliance and the SVC perform a rekey, they renegotiate the crypto keys and initialization vectors, increasing the security of the connection.
  - Renegotiation Interval—Specifies the number of minutes from the start of the session until the rekey takes place, from 1 through 10080 (1 week).
  - Renegotiation Method—Specifies whether and how SVC establishes a new tunnel during SVC rekey. If you check none, SVC rekey is disabled. If you check SSL, SSL renegotiation takes place during SVC rekey. If you select New tunnel, SVC establishes a new tunnel during SVC rekey. We recommend that you configure SSL as the rekey method.

- Dead Peer Detection—Dead Peer Detection (DPD) ensures that the security appliance (gateway) or the SVC can quickly detect a condition where the peer is not responding, and the connection has failed.
  - Gateway Side Detection—Enables DPD performed by the security appliance (gateway) and specifies the frequency, from 30 to 3600 seconds, with which the security appliance performs DPD. If you uncheck enable, DPD performed by the security appliance is disabled.
  - Client Side Detection—Enables DPD performed by the SVC (client), and specifies the frequency, from 30 to 3600 seconds, with which the SVC performs DPD. If you uncheck enable, DPD performed by the SVC is disabled.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Group Policy > WebVPN Tab > Auto Signon Tab

You can get to this panel through various paths.

The Auto Signon window or tab lets you configure or edit auto signon for WebVPN users. Auto signon is a simplified single signon method that you can use if you do not already have an SSO method deployed on your internal network. With auto signon configured for particular internal servers, the security appliance passes the login credentials that the WebVPN user used to login to the security appliance (username and password) to those particular internal servers. You configure the security appliance to respond to a specific authentication method for a particular range of servers. The authentication methods you can configure the security appliance to respond to are NTLM authentication, HTTP Basic authentication, or both methods.

Auto signon is a straight-forward method for configuring SSO for particular internal servers. This section describes the procedure for setting up SSO with auto signon. If you already have SSO deployed using Computer Associates' SiteMinder SSO server and want to configure the security appliance to support this solution, see SSO Servers. If you use SSO with HTTP Forms protocol and want to configure the security appliance to support this method, see AAA Setup.

Fields

- Inherit—Click to uncheck and allow WebVPN login credentials to be used to login to specific internal servers.
- IP Address—*Display only.* In conjunction with the following Mask, displays the IP address range of the servers to be authenticated to as configured with the Add/Edit Auto Signon dialog box. You can specify a server using either the server URI or the server IP address and mask.
- Mask—*Display only.* In conjunction with the preceding IP Address, displays the IP address range of the servers configured to support auto signon with the Add/Edit Auto Signon dialog box.

- URI—*Display only.* Displays a URI mask that identifies the servers configured with the Add/Edit Auto Signon dialog box.
- Authentication Type—*Display only.* Displays the type of authentication—basic HTTP, NTLM, or basic and NTLM—as configured with the Add/Edit Auto Signon dialog box.
- Add/Edit—Click to add or edit an auto signon instruction. An auto signon instruction defines a range of internal servers using the auto signon feature and the particular authentication method.
- Delete—Click to delete an auto signon instruction selected in the Auto Signon table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

ACLs

Configuration > VPN > Web VPN > ACLs

This window lets you configure ACLs for WebVPN.

Fields

- View (Unlabeled)—Indicates whether the selected entry is expanded (minus sign) or contracted (plus sign).
- # column—Specifies the ACE ID number.
- Enable—Indicates whether this ACL is enabled or disabled. You can enable or disable the ACL using this check box.
- Action—Specifies whether this ACL permits or denies access.
- Type—Specifies whether this ACL applies to a URL or a TCP address/port.
- Filter—Specifies the type of filter being applied.
- Syslog Level (Interval)—Specifies the syslog parameters for this ACL.
- Time Range—Specifies the name of the time range, if any, for this ACL. The time range can be a single interval or a series of periodic ranges.
- Description—Specifies the description, if any, of the ACL.
- Add ACL—Displays the Add Web Type ACL dialog box, in which you can specify an ACL ID.
- Add ACE—Displays the Add Web Type ACE dialog box, in which you specify parameters for the named ACL. This button is active only if there are one or more entries in the Web Type ACL table.
- Edit ACE/Delete—Click to edit or delete the highlighted ACL or ACE. When you delete an ACL, you also delete all of its ACEs. No warning or undelete.
- Move Up/Move Down—Highlight an ACL or ACE and click these buttons to change the order of ACLs and ACEs. The security appliance checks WebVPN ACLs and their ACEs in priority order according to their position in the ACLs list box until it finds a match.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Tunnel Group

Configuration > VPN > General > Tunnel Group

The parameters in the Tunnel Group window let you manage VPN tunnel groups. A VPN tunnel group represents a connection-specific record for IPSec and WebVPN connections.

The IPSec group uses the IPSec tunnel-group parameters to create the tunnel. An IPSec tunnel group can be either remote-access or LAN-to-LAN. The IPSec group is configured on the internal server or on an external RADIUS server. For ASA 5505 in client mode or VPN 3002 hardware client parameters, which enable or disable interactive hardware client authentication and individual user authentication, the IPSec tunnel group parameters take precedence over parameters set for users and groups.

The WebVPN tunnel-group parameters are the parameters of the WebVPN group that you want to apply to this tunnel group. You configure WebVPN access on the Configuration > WebVPN window.

Fields

- Tunnel Group—Shows the configured parameters for existing VPN tunnel groups. The Tunnel Group table contains the following columns:
  - Name—Specifies the name or IP address of the tunnel group.
  - Type—Indicates the type of tunnel; for example, ipsec-l2l indicates an IPSec LAN-to-LAN tunnel. The other possibilities are ipsec-ra (IPSec remote access) and webvpn.
  - Group Policy—Indicates the name of the group policy for this tunnel group.
- Add—Offers a menu letting you choose a tunnel type: IPSec for Remote Access, IPSec for LAN-to-LAN Access, or WebVPN Access, and opens a dialog box on which you can configure the new tunnel group.
- Edit—Opens a dialog box that lets you modify an existing tunnel group.
- Delete—Removes the selected tunnel group from the list.
- Group Delimiter—Lets you select the delimiter character to use when parsing tunnel group names from the user names that the security appliance receives when tunnels are being negotiated. By default, no delimiter is specified, disabling group-name parsing.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Tunnel Group > General Tab > Basic Tab

You can get to this panel through various paths.

The Add or Edit window, General tab, Basic tab lets you specify a name for the tunnel group that you are adding, lets you select the group policy, and lets you specify whether to strip the realm and/or group from the username before passing it on to the AAA server. You can also configure password management.

On the Edit Tunnel Group window, the General tab displays the name and type of the selected tunnel group. All other functions are the same as for the Add Tunnel Group window.

Fields

- Name—Specifies the name assigned to this tunnel group. For the Edit function, this field is display-only.
- Type—Displays the type of tunnel group you are adding or editing. For Edit, this is a display-only field whose contents depend on your selection in the Add window.
- Group Policy—Lists the currently configured group policies. The default value is the default group policy, DfltGrpPolicy.
- Strip the realm (administrative domain) from the username before passing it on to the AAA server—Enables or disables stripping the realm from the username before passing the username on to the AAA server. Check the Strip Realm check box to remove the realm qualifier of the username during authentication. You can append the realm name to the username for AAA: authorization, authentication and accounting. The only valid delimiter for a realm is the @ character. The format is username@realm, for example, JaneDoe@it.cisco.com. If you check this Strip Realm check box, authentication is based on the username alone. Otherwise, authentication is based on the full username@realm string. You must check this box if your server is unable to parse delimiters.

Note You can append both the realm and the group to a username, in which case the security appliance uses parameters configured for the group *and* for the realm for AAA functions. The format for this option is *username[@realm]]<#or!>group]*, for example, *JaneDoe@it.cisco.com#VPNGroup*. If you choose this option, you must use either the # or ! character for the group delimiter because the security appliance cannot interpret the @ as a group delimiter if it is also present as the realm delimiter.

A Kerberos realm is a special case. The convention in naming a Kerberos realm is to capitalize the DNS domain name associated with the hosts in the Kerberos realm. For example, if users are in the it.cisco.com domain, you might call your Kerberos realm IT.CISCO.COM.

- Strip the group from the username before passing it on to the AAA server—Enables or disables stripping the group name from the username before passing the username on to the AAA server. Check Strip Group to remove the group name from the username during authentication. This option is meaningful only when you have also checked the Enable Group Lookup box. When you append

a group name to a username using a delimiter, and enable Group Lookup, the security appliance interprets all characters to the left of the delimiter as the username, and those to the right as the group name. Valid group delimiters are the @, #, and ! characters, with the @ character as the default for Group Lookup. You append the group to the username in the format *username<delimiter>group*, the possibilities being, for example, *JaneDoe@VPNGroup, JaneDoe#VPNGroup*, and *JaneDoe!VPNGroup*.

- Password Management—Lets you configure parameters relevant to overriding an account-disabled indication from a AAA server and to notifying users about password expiration.
  - Override account-disabled indication from AAA server—Overrides an account-disabled indication from a AAA server.

Note Allowing override account-disabled is a potential security risk.

  - Enable notification upon password expiration to allow user to change password—Checking this check box makes the following two parameters available. If you do not also check the Enable notification prior to expiration check box, the user receives notification only after the password has expired.
  - Enable notification prior to expiration—When you check this option, the security appliance notifies the remote user at login that the current password is about to expire or has expired, then offers the user the opportunity to change the password. If the current password has not yet expired, the user can still log in using that password. This parameter is valid for AAA servers that support such notification; that is, RADIUS, RADIUS with an NT server, and LDAP servers. The security appliance ignores this command if RADIUS or LDAP authentication has not been configured.

    Note that this does not change the number of days before the password expires, but rather, it enables the notification. If you check this check box, you must also specify the number of days.
  - Notify...days prior to expiration—Specifies the number of days before the current password expires to notify the user of the pending expiration. The range is 1 through 180 days.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Tunnel Group > General Tab > Authentication Tab

You can get to this panel through various paths.

This tab is available for IPSec on Remote Access and LAN-to-LAN tunnel groups. The settings on this tab apply to the tunnel group globally across the security appliance. To set authentication server group settings per interface, click the Advanced tab. The Add or Edit Tunnel Group window > General tab > Authentication tab lets you configure the following attributes:

- Authentication Server Group—Lists the available authentication server groups, including the LOCAL group (the default). You can also select None. Selecting something other than None or Local makes available the Use LOCAL if Server Group Fails check box. To set the authentication server group per interface, go to the Advanced tab.
- Use LOCAL if Server Group fails—Enables or disables fallback to the LOCAL database if the group specified by the Authentication Server Group attribute fails.
- NAC Authentication Server Group—Specifies the authentication server group to use for posture validation. This field is active only if you have configured NAC on the security appliance. You must have an ACS group consisting of at least one server configured to support NAC. The list displays the names of all server groups of type RADIUS configured on this security appliance that are available for remote access tunnels.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Tunnel Group > General Tab > Authorization Tab

You can get to this panel through various paths.

This tab is available for IPSec on Remote Access and LAN-to-LAN tunnel groups. The settings on this tab apply to the tunnel group globally across the security appliance. On this tab, you can configure the following attributes:

- Authorization Server Group—Lists the available authorization server groups, including the LOCAL group. When VPN Authorization is defined as LOCAL, the attributes configured in the default group policy DfltGrpPolicy are enforced. You can also select None (the default). Selecting something other than None makes available the check box for Users must exist in authorization database to connect.
- Users must exist in the authorization database to connect—Tells the security appliance to allow only users in the authorization database to connect. By default this feature is disabled. You must have a configured authorization server to use this feature.
- Interface-Specific Authorization Server Groups—(Optional) Lets you configure authorization server groups on a per-interface basis. Interface-specific authorization server groups take precedence over the global server group. If you do not explicitly configure interface-specific authorization, authorization takes place only at the group level.
  - Interface—Select the interface on which to perform authorization. The standard interfaces are outside (the default), inside, and DMZ. If you have configured other interfaces, they also appear in the list.
  - Server Group—Select an available, previously configured authorization server group or group of servers, including the LOCAL group. You can associate a server group with more than one interface.
  - Add—Click Add to add the interface/server group setting to the table and remove the interface from the available list.

  - Remove—Click Remove to remove the interface/server group from the table and restore the interface to the available list.
- Authorization Settings—Lets you set values for usernames that the security appliance recognizes for authorization. This applies to users that authenticate with digital certificates and require LDAP or RADIUS authorization.
  - Use the entire DN as the username—Allows the use of the entire Distinguished Name (DN) as the username.
  - Specify individual DN fields as the username—Enables the use of individual DN fields as the username.
  - Primary DN Field—Lists all of the DN field identifiers for your selection.

DN Field	Definition
Country (C)	Two-letter country abbreviation. These codes conform to ISO 3166 country abbreviations.
Common Name (CN)	Name of a person, system, or other entity. This is the lowest (most specific) level in the identification hierarchy.
DN Qualifier (DNQ)	Specific DN attribute.
E-mail Address (EA)	E-mail address of the person, system or entity that owns the certificate.
Generational Qualifier (GENQ)	Generational qualifier such as Jr., Sr., or III.
Given Name (GN)	First name of the certificate owner.
Initials (I)	First letters of each part of the certificate owner's name.
Locality (L)	City or town where the organization is located.
Name (N)	Name of the certificate owner.
Organization (O)	Name of the company, institution, agency, association, or other entity.
Organizational Unit (OU)	Subgroup within the organization.
Serial Number (SER)	Serial number of the certificate.
Surname (SN)	Family name or last name of the certificate owner.
State/Province (S/P)	State or province where the organization is located.
Title (T)	Title of the certificate owner, such as Dr.
User ID (UID)	Identification number of the certificate owner.
User Principal Name (UPN)	Used with Smart Card certificate authentication.

  - Secondary DN Field—Lists all of the DN field identifiers (see the foregoing table) for your selection and adds the option None for no selection.

Add/Edit Tunnel Group > General Tab > Accounting Tab

You can get to this panel through various paths.

This tab is available for IPSec on Remote Access and LAN-to-LAN tunnel groups. The setting on this tab applies to the tunnel group globally across the security appliance. The Add or Edit Tunnel Group window > General tab > Accounting tab lets you configure the following attribute:

- Accounting Server Group—Lists the available accounting server groups. You can also select None (the default). LOCAL is not an option.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Tunnel Group > General Tab > Client Address Assignment Tab

You can get to this panel through various paths.

To specify whether to use DHCP or address pools for address assignment, go to Configuration > VPN > I P Address Management > Assignment. The Add or Edit Tunnel Group window > General tab > Client Address Assignment tab, lets you configure the following Client Address Assignment attributes:

- DHCP Servers—Specifies a DHCP server to use. You can add up to 10 servers, one at a time.
  - IP Address—Specifies the IP address of a DHCP server.
  - Add—Adds the specified DHCP server to the list for client address assignment.
  - Delete—Deletes the specified DHCP server from the list for client address assignment. There is no confirmation or undo.
- Address Pools—Lets you specify up to 6 address pools, using the following parameters:
  - Available Pools—Lists the available, configured address pools you can choose.
  - Add—Adds the selected address pool to the list for client address assignment.
  - Remove—Moves the selected address pool from the Assigned Pools list to the Available Pools list.
  - Assigned Pools—Lists the address pools selected for address assignment.

Note To configure interface-specific address pools, go to the Advanced tab.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Tunnel Group > General Tab > Advanced Tab

You can get to this panel through various paths.

The Add or Edit Tunnel Group window, General tab, Advanced tab, lets you configure the following interface-specific attributes:

- Interface-Specific Authentication Server Groups—Lets you configure an interface and server group for authentication.
  - Interface—Lists available interfaces for selection.
  - Server Group—Lists authentication server groups available for this interface.
  - Use LOCAL if server group fails—Enables or disables fallback to the LOCAL database if the server group fails.
  - Add—Adds the association between the selected available interface and the authentication server group to the assigned list.
  - Remove—Moves the selected interface and authentication server group association from the assigned list to the available list.
  - Interface/Server Group/Use Fallback—Show the selections you have added to the assigned list.
- Interface-Specific Client IP Address Pools—-Lets you specify an interface and Client IP address pool. You can have up to 6 pools.
  - Interface—Lists the available interfaces to add.
  - Address Pool—Lists address pools available to associate with this interface.
  - Add—Adds the association between the selected available interface and the client IP address pool to the assigned list.
  - Remove—Moves the selected interface/address pool association from the assigned list to the available list.
  - Interface/Address Pool—Shows the selections you have added to the assigned list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Tunnel Group > IPSec for Remote Access > IPSec Tab

Configuration > VPN > General > Tunnel Group > Add/Edit Tunnel Group > IPSec for Remote Access > IPSec Tab

On the Add or Edit Tunnel Group window for IPSec for Remote Access, the IPSec tab lets you configure or edit IPSec-specific tunnel group parameters.

Fields

- Pre-shared Key—Lets you specify the value of the pre-shared key for the tunnel group. The maximum length of the pre-shared key is 128 characters.

- Trustpoint Name—Selects a trustpoint name, if any trustpoints are configured. A trustpoint is a representation of a certificate authority. A trustpoint contains the identity of the CA, CA-specific configuration parameters, and an association with one enrolled identity certificate.
- Authentication Mode—Specifies the authentication mode: none, xauth, or hybrid.
  - none—Specifies no authentication mode.
  - xauth—Specifies the use of IKE Extended Authentication mode, which provides the capability of authenticating a user within IKE using TACACS+ or RADIUS.
  - hybrid—Specifies the use of Hybrid mode, which lets you use digital certificates for security appliance authentication and a different, legacy method—such as RADIUS, TACACS+ or SecurID—for remote VPN user authentication. This mode breaks phase 1 of the Internet Key Exchange (IKE) into the following steps, together called hybrid authentication:

  1. The security appliance authenticates to the remote VPN user with standard public key techniques. This establishes an IKE security association that is unidirectionally authenticated.
  2. An extended authentication (xauth) exchange then authenticates the remote VPN user. This extended authentication can use one of the supported legacy authentication methods.

Note Before setting the authentication type to hybrid, you must configure the authentication server and create a pre-shared key.

- IKE Peer ID Validation—Selects whether IKE peer ID validation is ignored, required, or checked only if supported by a certificate.
- Enable sending certificate chain—Enables or disables sending the entire certificate chain. This action includes the root certificate and any subordinate CA certificates in the transmission.
- ISAKMP Keep Alive—Enables and configures ISAKMP keep alive monitoring.
  - Disable Keep Alives—Enables or disables ISAKMP keep alives.
  - Monitor Keep Alives—Enables or disables ISAKMP keep alive monitoring. Selecting this option makes available the Confidence Interval and Retry Interval fields.
  - Confidence Interval—Specifies the ISAKMP keep alive confidence interval. This is the number of seconds the security appliance should allow a peer to idle before beginning keepalive monitoring. The minimum is 10 seconds; the maximum is 300 seconds. The default for a remote access group is 300 seconds.
  - Retry Interval—Specifies number of seconds to wait between ISAKMP keep alive retries. The default is 2 seconds.
  - Head end will never initiate keepalive monitoring—Specifies that the central-site security appliance never initiates keepalive monitoring.
- Interface-Specific Authentication Mode—Specifies the authentication mode on a per-interface basis.
  - Interface—Lets you select the interface name. The default interfaces are inside and outside, but if you have configured a different interface name, that name also appears in the list.
  - Authentication Mode—Lets you select the authentication mode, none, xauth, or hybrid, as above.
  - Interface/Authentication Mode table—Shows the interface names and their associated authentication modes that are selected.

  - Add—Adds an interface/authentication mode pair selection to the Interface/Authentication Modes table.
  - Remove—Removes an interface/authentication mode pair selection from the Interface/Authentication Modes table.
- Client VPN Software Update Table—Lists the client type, VPN Client revisions, and image URL for each client VPN software package installed. For each client type, you can specify the acceptable client software revisions and the URL or IP address from which to download software upgrades, if necessary. The client update mechanism (described in detail under the Client Update window) uses this information to determine whether the software each VPN client is running is at an appropriate revision level and, if appropriate, to provide a notification message and an update mechanism to clients that are running outdated software.
  - Client Type—Identifies the VPN client type.
  - VPN Client Revisions—Specifies the acceptable revision level of the VPN client.
  - Image URL—Specifies the URL or IP address from which the correct VPN client software image can be downloaded. For Windows-based VPN clients, the URL must be of the form http:// or https://. For ASA 5505 in client mode or VPN 3002 hardware clients, the URL must be of the form tftp://.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Tunnel Group > PPP Tab

Configuration > VPN > General > Tunnel Group > Add/Edit Tunnel Group > PPP Tab

On the Add or Edit Tunnel Group window for a IPSec remote access tunnel group, the PPP tab lets you configure or edit the authentication protocols permitted of a PPP connection. This tab applies *only* to IPSec remote access tunnel groups.

Fields

- CHAP—Enables the use of the CHAP protocol for a PPP connection.
- MS-CHAP-V1—Enables the use of the MS-CHAP-V1 protocol for a PPP connection.
- MS-CHAP-V2—Enables the use of the MA-CHAP-V2 protocol for a PPP connection.
- PAP—Enables the use of the PAP protocol for a PPP connection.
- EAP-PROXY—Enables the use of the EAP-PROXY protocol for a PPP connection. EAP refers to the Extensible Authentication protocol.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Tunnel Group > IPSec for LAN to LAN Access > General Tab > Basic Tab

Configuration > VPN > General > Tunnel Group > Add/Edit Tunnel Group > IPSec for LAN to LAN Access > General Tab > Basic Tab

On the Add or Edit Tunnel Group window for LAN-to-LAN Remote Access, the General tab, Basic tab you can specify a name for the tunnel group that you are adding (Add function only) and select the group policy.

On the Edit Tunnel Group window, the General tab displays the name and type of the tunnel group you are modifying.

Fields

- Name—Specifies the name assigned to this tunnel group. For the Edit function, this field is display-only.
- Type—(*Display-only*) Displays the type of tunnel group you are adding or editing. The contents of this field depend on your selection on the previous window.
- Group Policy—Lists the currently configured group policies. The default value is the default group policy, DfltGrpPolicy.
- Strip the realm (administrative domain) from the username before passing it on to the AAA server—Enables or disables stripping the realm from the username before passing the username on to the AAA server. Check the Strip Realm check box to remove the realm qualifier of the username during authentication. You can append the realm name to the username for AAA: authorization, authentication and accounting. The only valid delimiter for a realm is the @ character. The format is username@realm, for example, JaneDoe@it.cisco.com. If you check this Strip Realm check box, authentication is based on the username alone. Otherwise, authentication is based on the full username@realm string. You must check this box if your server is unable to parse delimiters.

Note You can append both the realm and the group to a username, in which case the security appliance uses parameters configured for the group *and* for the realm for AAA functions. The format for this option is *username[@realm]]<#or!>group]*, for example, *JaneDoe@it.cisco.com#VPNGroup*. If you choose this option, you must use either the # or ! character for the group delimiter because the security appliance cannot interpret the @ as a group delimiter if it is also present as the realm delimiter.

A Kerberos realm is a special case. The convention in naming a Kerberos realm is to capitalize the DNS domain name associated with the hosts in the Kerberos realm. For example, if users are in the it.cisco.com domain, you might call your Kerberos realm IT.CISCO.COM.

- Strip the group from the username before passing it on to the AAA server—Enables or disables stripping the group name from the username before passing the username on to the AAA server. Check Strip Group to remove the group name from the username during authentication. This option is meaningful only when you have also checked the Enable Group Lookup box. When you append

a group name to a username using a delimiter, and enable Group Lookup, the security appliance interprets all characters to the left of the delimiter as the username, and those to the right as the group name. Valid group delimiters are the @, #, and ! characters, with the @ character as the default for Group Lookup. You append the group to the username in the format *username<delimiter>group*, the possibilities being, for example, *JaneDoe@VPNGroup, JaneDoe#VPNGroup*, and *JaneDoe!VPNGroup*.

- Password Management—Lets you configure parameters relevant to overriding an account-disabled indication from a AAA server and to notifying users about password expiration.
  - Override account-disabled indication from AAA server—Overrides an account-disabled indication from a AAA server.

Note Allowing override account-disabled is a potential security risk.

  - Enable notification upon password expiration to allow user to change password—Checking this check box makes the following two parameters available. If you do not also check the Enable notification prior to expiration check box, the user receives notification only after the password has expired.
  - Enable notification prior to expiration—When you check this option, the security appliance notifies the remote user at login that the current password is about to expire or has expired, then offers the user the opportunity to change the password. If the current password has not yet expired, the user can still log in using that password. This parameter is valid for AAA servers that support such notification; that is, RADIUS, RADIUS with an NT server, and LDAP servers. The security appliance ignores this command if RADIUS or LDAP authentication has not been configured.

    Note that this does not change the number of days before the password expires, but rather, it enables the notification. If you check this check box, you must also specify the number of days.
  - Notify...days prior to expiration—Specifies the number of days before the current password expires to notify the user of the pending expiration. The range is 1 through 180 days.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Tunnel Group > IPSec for LAN to LAN Access > IPSec Tab

Configuration > VPN > General > Tunnel Group > Add/Edit Tunnel Group > IPSec for LAN to LAN Access > IPSec Tab

The Add or Edit Tunnel Group window for IPSec for LAN-to-LAN access, IPSec tab, lets you configure or edit IPSec LAN-to-LAN-specific tunnel group parameters.

Fields

- Name—Specifies the name assigned to this tunnel group. For the Edit function, this field is display-only.
- Type—(*Display-only*) Displays the type of tunnel group you are adding or editing. The contents of this field depend on your selection on the previous window.
- Pre-shared Key—Lets you specify the value of the pre-shared key for the tunnel group. The maximum length of the pre-shared key is 128 characters.
- Trustpoint Name—Selects a trustpoint name, if any trustpoints are configured. A trustpoint is a representation of a certificate authority. A trustpoint contains the identity of the CA, CA-specific configuration parameters, and an association with one enrolled identity certificate.
- Authentication Mode—Specifies the authentication mode: none, xauth, or hybrid.
  - none—Specifies no authentication mode.
  - xauth—Specifies the use of IKE Extended Authentication mode, which provides the capability of authenticating a user within IKE using TACACS+ or RADIUS.
  - hybrid—Specifies the use of Hybrid mode, which lets you use digital certificates for security appliance authentication and a different, legacy method—such as RADIUS, TACACS+ or SecurID—for remote VPN user authentication. This mode breaks phase 1 of the Internet Key Exchange (IKE) into the following steps, together called hybrid authentication:

  1. The security appliance authenticates to the remote VPN user with standard public key techniques. This establishes an IKE security association that is unidirectionally authenticated.
  2. An extended authentication (xauth) exchange then authenticates the remote VPN user. This extended authentication can use one of the supported legacy authentication methods.

Note Before setting the authentication type to hybrid, you must configure the authentication server and create a pre-shared key.

- IKE Peer ID Validation—Selects whether IKE peer ID validation is ignored, required, or checked only if supported by a certificate.
- Enable sending certificate chain—Enables or disables sending the entire certificate chain. This action includes the root certificate and any subordinate CA certificates in the transmission.
- ISAKMP Keep Alive—Enables and configures ISAKMP keep alive monitoring.
  - Disable Keep Alives—Enables or disables ISAKMP keep alives.
  - Monitor Keep Alives—Enables or disables ISAKMP keep alive monitoring. Selecting this option makes available the Confidence Interval and Retry Interval fields.
  - Confidence Interval—Specifies the ISAKMP keep alive confidence interval. This is the number of seconds the security appliance should allow a peer to idle before beginning keepalive monitoring. The minimum is 10 seconds; the maximum is 300 seconds. The default for a remote access group is 300 seconds.
  - Retry Interval—Specifies number of seconds to wait between ISAKMP keep alive retries. The default is 2 seconds.
  - Head end will never initiate keepalive monitoring—Specifies that the central-site security appliance never initiates keepalive monitoring.
- Interface-Specific Authentication Mode—Specifies the authentication mode on a per-interface basis.

  - Interface—Lets you select the interface name. The default interfaces are inside and outside, but if you have configured a different interface name, that name also appears in the list.
  - Authentication Mode—Lets you select the authentication mode, none, xauth, or hybrid, as above.
  - Interface/Authentication Mode table—Shows the interface names and their associated authentication modes that are selected.
  - Add—Adds an interface/authentication mode pair selection to the Interface/Authentication Modes table.
  - Remove—Removes an interface/authentication mode pair selection from the Interface/Authentication Modes table.
- Client VPN Software Update Table—Lists the client type, VPN Client revisions, and image URL for each client VPN software package installed. For each client type, you can specify the acceptable client software revisions and the URL or IP address from which to download software upgrades, if necessary. The client update mechanism (described in detail under the Client Update window) uses this information to determine whether the software each VPN client is running is at an appropriate revision level and, if appropriate, to provide a notification message and an update mechanism to clients that are running outdated software.
  - Client Type—Identifies the VPN client type.
  - VPN Client Revisions—Specifies the acceptable revision level of the VPN client.
  - Image URL—Specifies the URL or IP address from which the correct VPN client software image can be downloaded. For Windows-based VPN clients, the URL must be of the form http:// or https://. For ASA 5505 in client mode or VPN 3002 hardware clients, the URL must be of the form tftp://.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Tunnel Group > WebVPN Access > General Tab > Basic Tab

Configuration > VPN > General > Tunnel Group > Add/Edit> WebVPN Access > General Tab > Basic Tab

The Add or Edit pane, General tab, Basic tab lets you specify a name for the tunnel group that you are adding, lets you select the group policy, and lets you configure password management.

On the Edit Tunnel Group window, the General tab displays the name and type of the selected tunnel group. All other functions are the same as for the Add Tunnel Group window.

Fields

- Name—Specifies the name assigned to this tunnel group. For the Edit function, this field is display-only.

- Type—Displays the type of tunnel group you are adding or editing. For Edit, this is a display-only field whose contents depend on your selection in the Add window.
- Group Policy—Lists the currently configured group policies. The default value is the default group policy, DfltGrpPolicy.
- Strip the realm —Not available for WebVPN.
- Strip the group —Not available or WebVPN.
- Password Management—Lets you configure parameters relevant to overriding an account-disabled indication from a AAA server and to notifying users about password expiration.
  - Override account-disabled indication from AAA server—Overrides an account-disabled indication from a AAA server.

Note Allowing override account-disabled is a potential security risk.

  - Enable notification upon password expiration to allow user to change password—Checking this check box makes the following two parameters available. If you do not also check the Enable notification prior to expiration check box, the user receives notification only after the password has expired.
  - Enable notification prior to expiration—When you check this option, the security appliance notifies the remote user at login that the current password is about to expire or has expired, then offers the user the opportunity to change the password. If the current password has not yet expired, the user can still log in using that password. This parameter is valid for AAA servers that support such notification; that is, RADIUS, RADIUS with an NT server, and LDAP servers. The security appliance ignores this command if RADIUS or LDAP authentication has not been configured.

    Note that this does not change the number of days before the password expires, but rather, it enables the notification. If you check this check box, you must also specify the number of days.
  - Notify...days prior to expiration—Specifies the number of days before the current password expires to notify the user of the pending expiration. The range is 1 through 180 days.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Tunnel Group > WebVPN Tab > Basic Tab

Configuration > VPN > General > Tunnel Group > Add/Edit> WebVPN Access > WebVPN Tab > Basic Tab

The attributes on the Add/Edit Tunnel Group General Tab tabs for WebVPN are the same as those for Add/Edit Tunnel Group General Tab tabs for IPSec Remote Access. The following description applies to the fields appearing on the WebVPN Tab tabs.

Fields

The Basic tab lets you configure the following WebVPN attributes:

- Authentication—Specifies the type of authentication to perform: AAA, Certificate, or Both. The default value is AAA.
- DNS Group—Specifies the DNS server to use for a WebVPN tunnel-group. The default value is DefaultDNS.
- CSD Failure group policy—This attribute is valid only for security appliances with Cisco Secure Desktop installed. The security appliance uses this attribute to limit access rights to remote CSD clients if you use Cisco Secure Desktop Manager to set the VPN feature policy to one of the following options:
  - "Use Failure Group-Policy."
  - "Use Success Group-Policy, if criteria match," and the criteria fail to match.

  This attribute specifies the name of the failure group policy to be applied. Choose a group policy to differentiate access rights from those associated with the default group policy. The default value is DfltGrpPolicy.

Note The security appliance does not use this attribute if you set the VPN feature policy to "Always use Success Group-Policy."

For more information, see the *Cisco Secure Desktop Configuration Guide for Cisco ASA 5500 Series Administration Guide*

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Tunnel Group > WebVPN Access > WebVPN Tab > NetBIOS Servers Tab

Configuration > VPN > General > Tunnel Group > Add/Edit Tunnel Group > WebVPN Access > WebVPN Tab > NetBIOS Servers Tab

The table on this tab shows the attributes of the already-configured NetBIOS servers. The Add or Edit Tunnel Group window for WebVPN Access, NetBIOS tab, lets you configure the NetBIOS attributes for the tunnel group. WebVPN uses NetBIOS and the Common Internet File System protocol to access or share files on remote systems. When you attempt a file-sharing connection to a Windows computer by using its computer name, the file server you specify corresponds to a specific NetBIOS name that identifies a resource on the network.

The security appliance queries NetBIOS name servers to map NetBIOS names to IP addresses. WebVPN requires NetBIOS to access or share files on remote systems.

To make the NBNS function operational, you must configure at least one NetBIOS server (host). You can configure up to 3 NBNS servers for redundancy. The security appliance uses the first server on the list for NetBIOS/CIFS name resolution. If the query fails, it uses the next server.

Fields

- IP Address—Displays the IP addresses of configured NetBIOS servers.
- Master Browser—Shows whether a server is a WINS server or one that can also be a CIFS server (that is, a master browser).
- Timeout (seconds)—Displays the initial time in seconds that the server waits for a response to an NBNS query before sending the query to the next server.
- Retries—Shows the number of times to retry sending an NBNS query to the configured servers, in order. In other words, this is the number of times to cycle through the list of servers before returning an error. The minimum number of retries is 0. The default number of retries is 2. The maximum number of retries is 10.
- Add/Edit—Click to add a NetBIOS server. This opens the Add or Edit NetBIOS Server dialog box.
- Delete—Removes the highlighted NetBIOS row from the list.
- Move Up/Move Down—The security appliance sends NBNS queries to the NetBIOS servers in the order in which they appear in this box. Use this box to change the priority order of the servers by moving them up or down in the list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Tunnel Group > WebVPN Access > WebVPN Tab > NetBIOS Servers Tab > Add/Edit NetBIOS Server

Configuration > VPN > General > Tunnel Group > Add/Edit Tunnel Group > WebVPN Access > WebVPN Tab > NetBIOS Servers Tab > Add/Edit NetBIOS Server

This dialog box lets you create a new entry for the NetBIOS servers table or modify an existing entry.

Fields

- IP Address—Specifies the IP address for the NetBIOS server.
- Master browser—Designates the current NetBIOS server as a master browser, rather than a WINS server.
- Timeout—Specifies the initial time in seconds the server waits for a response to an NBNS query before sending the query to the next server. The minimum time is 1 second. The default time is 2 seconds. The maximum time is 30 seconds. The time doubles with each retry cycle through the list of servers.
- Retries—Specifies the number of times to retry sending a NBNS query to the configured servers, in order. In other words, this is the number of times to cycle through the list of servers before returning an error. The minimum number of retries is 0. The default number of retries is 2. The maximum number of retries is 10.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Tunnel Group > WebVPN Access > WebVPN Tab > Group Aliases and URLs Tab

Configuration > VPN > General > Tunnel Group > Add/Edit Tunnel Group > WebVPN Access > WebVPN Tab > Group Aliases and URLs Tab

The Add or Edit Tunnel Group window for WebVPN Remote Access, Group Aliases and URLs tab, lets you specify alternative names for the group (group aliases) and specify incoming URLs for the group.

Specifying the group alias creates one or more alternate names by which the user can refer to a tunnel-group. The group alias that you specify here appears in the drop-down list on the login page. Each group can have multiple aliases or no alias. If you want the actual name of the tunnel group to appear on this list, specify it as an alias. This feature is useful when the same group is known by several common names, such as “Devtest” and “QA”.

Specifying a group URL eliminates the need for the user to select a group at login. When a user logs in, the security appliance looks for the user’s incoming URL in the tunnel-group-policy table. If it finds the URL and if this feature is enabled, then the security appliance automatically selects the appropriate server and presents the user with only the username and password fields in the login window. If the URL is disabled, then the dropdown list of groups is also displayed, and the user must make the selection.

You can configure multiple URLs (or no URLs) for a group. Each URL can be enabled or disabled individually. You must use a separate specification for each URL specified. You must specify the entire URL, which can use either the http or https protocol.

You cannot associate the same URL with multiple groups. The security appliance verifies the uniqueness of the URL before accepting it for a tunnel group.

Fields

- Group Aliases—Contains the following entries:
  - Alias—Specifies an alternative name for the tunnel group.
  - Add/Remove—Adds or removes a selected group alias from the list.
  - Enable—Enables the selected alias, so it appears on the dropdown list at logon. This check box is checked by default.

Note You cannot change the status of a disabled alias in the Alias/Status table merely by checking Enable and clicking OK, then Apply. You must first remove the disabled alias, then re-add it with the Enable check box checked.

  - Alias/Status —Shows whether each selected alias is enabled or disabled.
- Group URLs—Contains the following entries:
  - URL (http or https)—Specifies a URL to add to the list; for example, http://www.cisco.com. You must include the http:// or https:// protocol in the URL.

- Add/Remove—Adds or removes a selected group URL from the list.
- Enable—Enables the selected URL. The default is enabled.

Note You cannot change the status of a disabled URL in the URL/Status table merely by checking Enable and clicking OK, then Apply. You must first remove the disabled URL, then re-add it with the Enable check box checked.

- URL/Status—Shows whether each selected URL is enabled or disabled.

Example

You can set up different login screens for different groups by using a combination of customization profiles and tunnel groups. For example, assuming that you had created a customization profile called salesgui, you can create a WebVPN tunnel group called sales that refers to that customization profile, as the following example shows. This example displays the company logo instead of the default Cisco logo when the user logs in using WebVPN:

Step 1 Define a WebVPN customization named salesgui and change the default logo to mycompanylogo.gif. You must have previously loaded mycompanylogo.gif onto the flash memory of the security appliance and saved the configuration.

Step 2 Set up a username and associate it with the WebVPN customization you've just defined.

Step 3 Create a WebVPN tunnel-group named sales.

Step 4 Specify that you want to use the salesgui customization for this tunnel group.

Step 5 Set the group URL to the address that the user enters into the browser to log in to the security appliance; for example, if the security appliance has the IP address 192.168.3.3, set the group URL to https://192.168.3.3.

The security appliance maps this URL to the sales tunnel group and applies the salesgui customization profile to the login screen that the user sees.

Step 6 Save the configuration to memory.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Tunnel Group > WebVPN Access > WebVPN Tab > Web Page Tab

Configuration > VPN > General > Tunnel Group > Add/Edit Tunnel Group > WebVPN Access > WebVPN Tab > Web Page Tab

Use this tab to select a customized look and feel for the WebVPN end-user logon web page.

Fields

- Webpage Customization—Selects a previously defined web-page customization.
- New—Opens a dialog box in which you can configure a new web-page customization.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

For More Information

WebVPN End User Set-up

VPN System Options

Configuration > VPN > General > VPN System Options

The VPN System Options window lets you configure features specific to VPN sessions on the security appliance.

Fields

- Enable inbound IPSec sessions to bypass interface access-lists. Group policy and per-user authorization access lists still apply to the traffic—By default, the security appliance allows VPN traffic to terminate on a security appliance interface; you do not need to allow IKE or ESP (or other types of VPN packets) in an access rule. When this option is checked, you also do not need an access rule for local IP addresses of decrypted VPN packets. Because the VPN tunnel was terminated successfully using VPN security mechanisms, this feature simplifies configuration and maximizes the security appliance performance without any security risks. (Group policy and per-user authorization access lists still apply to the traffic.)

  You can require an access rule to apply to the local IP addresses by unchecking this option. The access rule applies to the local IP address, and not to the original client IP address used before the VPN packet was decrypted.
- Limit the maximum number of active IPSec VPN sessions—Enables or disables limiting the maximum number of active IPSec VPN sessions. The range depends on the hardware platform and the software license.
- Maximum Active IPSec VPN Sessions—Specifies the maximum number of active IPSec VPN sessions allowed. This field is active only when you select the preceding check box to limit the maximum number of active IPSec VPN sessions.
- L2TP Tunnel Keep-alive Timeout—Specifies the frequency, in seconds, of keepalive messages. The range is 10 through 300 seconds. The default is 60 seconds.
- Compression Settings—Specifies the features for which you want to enable compression: WebVPN, and SSL VPN Client. Compression is enabled by default.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Zone Labs Integrity Server

Configuration > VPN > General > Zone Labs Integrity Server

The Zone Labs Integrity Server panel lets you configure the security appliance to support a Zone Labs Integrity Server. This server is part of the Integrity System, a system designed to enforce security policies on remote clients entering the private network. In essence, the security appliance acts as a proxy for the client PC to the Firewall Server and relays all necessary Integrity information between the Integrity client and the Integrity server.

Note The current release of the security appliance supports one Integrity Server at a time even though the user interfaces support the configuration of up to five Integrity Servers. If the active Server fails, configure another Integrity Server on the security appliance and then reestablish the client VPN session.

Fields

- Server IP address—Type the IP address of the Integrity Server. Use dotted decimal notation.
- Add—Adds a new server IP address to the list of Integrity Servers. This button is active when an address is entered in the Server IP address field.
- Delete—Deletes the selected server from the list of Integrity Servers.
- Move Up—Moves the selected server up in the list of Integrity Servers. This button is available only when there is more than one server in the list.
- Move Down—Moves the selected server down in the list of Integrity Servers. This button is available only when there is more than one server in the list.
- Server Port—Type the security appliance port number on which it listens to the active Integrity server. This field is available only if there is at least one server in the list of Integrity Servers. The default port number is 5054, and it can range from 10 to 10000. This field is only available when there is a server in the Integrity Server list.
- Interface—Choose the interface security appliance interface on which it communicates with the active Integrity Server. This interface name menu is only available when there is a server in the Integrity Server list.
- Fail Timeout—Type the number of seconds that the security appliance should wait before it declares the active Integrity Server to be unreachable. The default is 10 and the range is from 5 to 20.
- Enable SSL Authentication—Check to enable authentication of the remote client SSL certificate by the security appliance. By default, client SSL authentication is disabled.
- Close connection on timeout—Check to close the connection between the security appliance and the Integrity Server on a timeout. By default, the connection remains open.

- Apply—Click to apply the Integrity Server setting to the security appliance running configuration.
- Reset—Click to remove Integrity Server configuration changes that have not yet been applied.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Easy VPN Remote

Configuration > VPN > Easy VPN Remote

Easy VPN Remote lets the ASA 5505 act as an Easy VPN client device. The ASA 5505 can then initiate a VPN tunnel to an Easy VPN server, which can be a security appliance, a Cisco VPN 3000 Concentrator, an IOS-based router, or a firewall acting as an Easy VPN server.

The Easy VPN client supports one of two modes of operation: Client Mode or Network Extension Mode (NEM). The mode of operation determines whether the Easy VPN Client inside hosts are accessible from the Enterprise network over the tunnel. Specifying a mode of operation is mandatory before making a connection because Easy VPN Client does not have a default mode.

Client mode, also called Port Address Translation (PAT) mode, isolates all devices on the Easy VPN Client private network from those on the enterprise network. The Easy VPN Client performs Port Address Translation (PAT) for all VPN traffic for its inside hosts. IP address management is neither required for the Easy VPN Client inside interface or the inside hosts.

NEM makes the inside interface and all inside hosts routable across the enterprise network over the tunnel. Hosts on the inside network obtain their IP addresses from an accessible subnet (statically or via DHCP) pre-configured with static IP addresses. PAT does not apply to VPN traffic in NEM. This mode does not require a VPN configuration for each client. The Cisco ASA 5505 configured for NEM mode supports automatic tunnel initiation. The configuration must store the group name, user name, and password. Automatic tunnel initiation is disabled if secure unit authentication is enabled.

The network and addresses on the private side of the Easy VPN Client are hidden, and cannot be accessed directly.

Fields

- Enable Easy VPN Remote—Enables the Easy VPN Remote feature and makes available the rest of the fields on this window for configuration.
- Mode—Selects either Client mode or Network extension mode.
  - Client mode—Uses Port Address Translation (PAT) mode to isolate the addresses of the inside hosts, relative to the client, from the enterprise network.
  - Network extension mode—Makes those addresses accessible from the enterprise network.

Note If the Easy VPN Remote is using NEM and has connections to secondary servers, establish an ASDM connection to each headend and check Enable Reverse Route Injection on the Configuration > VPN > IPSec > IPSec Rules > Tunnel Policy (Crypto Map) - Advanced tab to configure dynamic announcements of the remote network using RRI.

  - Auto connect—The Easy VPN Remote establishes automatic IPSec data tunnels unless both of the following are true: Network extension mode is configured locally, and split-tunneling is configured on the group policy pushed to the Easy VPN Remote. If both are true, checking this attribute automates the establishment of IPSec data tunnels. Otherwise, this attribute has no effect.

- Group Settings—Specifies whether to use a pre-shared key or an X.509 certificate for user authentication.
  - Pre-shared key—Enables the use of a pre-shared key for authentication and makes available the subsequent Group Name, Group Password, and Confirm Password fields for specifying the group policy name and password containing that key.
  - Group Name—Specifies the name of the group policy to use for authentication.
  - Group Password—Specifies the password to use with the specified group policy.
  - Confirm Password—Requires you to confirm the group password just entered.
  - X.509 Certificate—Specifies the use of an X.509 digital certificate, supplied by a Certificate Authority, for authentication.
  - Select Trustpoint—Lets you select a trustpoint, which can be an IP address or a hostname, from the drop-down list. To define a trustpoint, click the link to Trustpoint(s) configuration at the bottom of this area.
  - Send certificate chain—Enables sending a certificate chain, not just the certificate itself. This action includes the root certificate and any subordinate CA certificates in the transmission.
- User Settings—Configures user login information.
  - User Name—Configures the VPN username for the Easy VPN Remote connection. Xauth provides the capability of authenticating a user within IKE using TACACS+ or RADIUS. Xauth authenticates a user (in this case, the Easy VPN hardware client) using RADIUS or any of the other supported user authentication protocols. The Xauth username and password parameters are used when secure unit authentication is disabled and the server requests Xauth credentials. If secure unit authentication is enabled, these parameters are ignored, and the security appliance prompts the user for a username and password.
  - User Password—Configures the VPN user password for the Easy VPN Remote connection.
  - Confirm Password—Requires you to confirm the user password just entered.
- Easy VPN Server To Be Added—Adds or removes an Easy VPN server. Any ASA or VPN 3000 Concentrator Series can act as a Easy VPN server. A server must be configured before a connection can be established. The security appliance supports IPv4 addresses, the names database, or DNS names and resolves addresses in that order. The first server in the Easy VPN Server(s) list is the primary server. You can specify a maximum of ten backup servers in addition to the primary server.
  - Name or IP Address—The name or IP address of an Easy VPN server to add to the list.
  - Add—Moves the specified server to the Easy VPN Server(s) list.

- Remove—Moves the selected server from the Easy VPN Server(s) list to the Name or IP Address file. Once you do this, however, you cannot re-add the same address unless you re-enter the address in the Name or IP Address field.
- Easy VPN Server(s)—Lists the configured Easy VPN servers in priority order.
- Move Up/Move Down—Changes the position of a server in the Easy VPN Server(s) list. These buttons are available only when there is more than one server in the list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Advanced Easy VPN Properties

Configuration > VPN > Easy VPN Remote > Advanced

Device Pass-Through

Certain devices like Cisco IP phones, printers, and the like are incapable of performing authentication, and therefore of participating in individual unit authentication. To accommodate these devices, the device pass-through feature, enabled by the MAC Exemption attributes, exempts devices with the specified MAC addresses from authentication when Individual User Authentication is enabled.

The first 24 bits of the MAC address indicate the manufacturer of the piece of equipment. The last 24 bits are the unit's serial number in hexadecimal format.

Tunneled Management

When operating an ASA model 5505 device behind a NAT device, use the Tunneled Management attributes to specify how to configure device management— in the clear or through the tunnel—and specify the network or networks allowed to manage the Easy VPN Remote connection through the tunnel. The public address of the ASA 5505 is not accessible when behind the NAT device unless you add static NAT mappings on the NAT device.

When operating a Cisco ASA 5505 behind a NAT device, use the **vpnclient management** command to specify how to configure device management— with additional encryption or without it—and specify the hosts or networks to be granted administrative access. The public address of the ASA 5505 is not accessible when behind the NAT device unless you add static NAT mappings on the NAT device.

Fields

- MAC Exemption—Configures a set of MAC addresses and masks used for device pass-through for the Easy VPN Remote connection
  - MAC Address—Exempts the device with the specified MAC address from authentication. The format for specifying the MAC address this field uses three hex digits, separated by periods; for example, 45ab.ff36.9999.

  - MAC Mask—The format for specifying the MAC mask in this field uses three hex digits, separated by periods; for example, the MAC mask ffff.ffff.ffff matches just the specified MAC address. A MAC mask of all zeroes matches no MAC address, and a MAC mask of ffff.ff00.0000 matches all devices made by the same manufacturer.
  - Add—Adds the specified MAC address and mask pair to the MAC Address/Mask list.
  - Remove—Moves the selected MAC address and mask pair from the MAC Address/MAC list to the individual MAC Address and MAC Mask fields.
- Tunneled Management—Configures IPSec encryption for device management and specifies the network or networks allowed to manage the Easy VPN hardware client connection through the tunnel. Selecting Clear Tunneled Management merely removes that IPSec encryption level and does not affect any other encryption, such as SSH or https, that exists on the connection.
  - Enable Tunneled Management—Adds a layer of IPSec encryption to the SSH or HTTPS encryption already present in the management tunnel.
  - Clear Tunneled Management—Uses the encryption already present in the management tunnel, without additional encryption.
  - IP Address— Specifies the IP address of the host or network to which you want to grant administrative access to the Easy VPN hardware client through the VPN tunnel. You can individually add one or more IP addresses and their respective network masks.
  - Mask—Specifies the network mask for the corresponding IP address.
  - Add—Moves the specified IP address and mask to the IP Address/Mask list.
  - Remove—Moves the selected IP address and mask pair from the IP Address/Mask list to the individual IP Address and Mask fields in this area.
  - IP Address/Mask—Lists the configured IP address and mask pairs to be operated on by the Enable or Clear functions in this area.
- IPSec Over TCP—Configure the Easy VPN Remote connection to use TCP-encapsulated IPSec.
  - Enable—Enables IPSec over TCP.

Note Choose Configuration > VPN > IPSec > Pre-Fragmentation, double-click the outside interface, and set the DF Bit Setting Policy to Clear if you configure the Easy VPN Remote connection to use TCP-encapsulated IPSec. The Clear setting lets the security appliance send large packets.

  - Enter Port Number—Specifies the port number to use for the IPSec over TCP connection.
- Server Certificate—Configures the Easy VPN Remote connection to accept only connections to Easy VPN servers with the specific certificates specified by the certificate map. Use this parameter to enable Easy VPN server certificate filtering. To define a certificate map, go to Configuration > VPN > IKE > Certificate Group Matching > Rules.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

CHAPTER 29

WebVPN

WebVPN lets users establish a secure, remote-access VPN tunnel to the security appliance using a web browser. There is no need for either a software or hardware client. WebVPN provides easy access to a broad range of web resources and both web-enabled and legacy applications from almost any computer that can reach HTTPS Internet sites. WebVPN uses Secure Socket Layer Protocol and its successor, Transport Layer Security (SSL/TLS1) to provide a secure connection between remote users and specific, supported internal resources that you configure at a central site. The security appliance recognizes connections that need to be proxied, and the HTTP server interacts with the authentication subsystem to authenticate users.

The network administrator provides access to WebVPN resources on a user or group basis. Users have no direct access to resources on the internal network.

WebVPN works on the platform in single, routed mode.

For information on configuring WebVPN for end users, see WebVPN End User Set-up.

WebVPN Security Precautions

WebVPN connections on the security appliance are very different from remote access IPSec connections, particularly with respect to how they interact with SSL-enabled servers, and precautions to reduce security risks.

In a WebVPN connection, the security appliance acts as a proxy between the end user web browser and target web servers. When a WebVPN user connects to an SSL-enabled web server, the security appliance establishes a secure connection and validates the server SSL certificate.

The current implementation of WebVPN does not permit communication with sites that present expired certificates. Nor does the security appliance perform trusted CA certificate validation. Therefore, WebVPN users cannot analyze the certificate an SSL-enabled web server presents before communicating with it.

To minimize the risks involved with SSL certificates:

- Configure a group policy for all users who need WebVPN access and enable the WebVPN feature only for that group policy.
- Limit Internet access for WebVPN users. One way to do this is to clear the **Enable URL entry** check box on the Configuration > VPN > General > Group Policy > WebVPN panel. Then configure links to specific targets within the private network (Configuration > VPN > WebVPN > Servers and URLs).

- Educate users. If an SSL-enabled site is not inside the private network, users should not visit this site over a WebVPN connection. They should open a separate browser window to visit such sites, and use that browser to view the presented certificate.

ACLs

Configuration > VPN > WebVPN > ACLs

You can configure ACLs (Access Control Lists) to apply to user sessions. These are filters that permit or deny user access to specific networks, subnets, hosts, and web servers.

- If you do not define any filters, all connections are permitted.
- The security appliance supports only an inbound ACL on an interface.
- At the end of each ACL, there is an implicit, unwritten rule that denies all traffic that is not permitted. If traffic is not explicitly permitted by an access control entry (ACE), the security appliance denies it. ACEs are referred to as rules in this topic.

This pane lets you add and edit WebVPN ACLs and the ACL entries that each ACL contains. It also displays summary information about ACLS and ACEs, and lets you enable or disable them, and change their priority order.

Fields

- Add ACL—Click to add an ACL or ACE. To insert a new ACE before or after an existing ACE, click Insert or Insert After.
- Edit—Click to edit the highlighted ACE. When you delete an ACL, you also delete all of its ACEs. No warning or undelete.
- Delete—Click to delete the highlighted ACL or ACE. When you delete an ACL, you also delete all of its ACEs. No warning or undelete.
- Move UP/Move Down—Highlight an ACL or ACE and click these buttons to change the order of ACLs and ACEs. The security appliance checks WebVPN ACLs and their ACEs in priority order according to their position in the ACLs list box until it finds a match.
- +/-—Click to expand (**+**) or collapse (**-**) to view or hide the list of ACEs under each ACL.
- No—Displays the priority of the ACEs under each ACL. The order in the list determines priority.
- Address—Displays the IP address or URL of the application or service to which the ACE applies.
- Service—Displays the TCP service to which the ACE applies.
- Action—Displays whether the ACE permits or denies WebVPN access.
- Time—Displays the time range associated with the ACE.
- Logging (Interval)—Displays the configured logging behavior, either disabled or with a specified level and time interval.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add ACL

This pane lets you create a new ACL.

Fields

- ACL Name—Enter a name for the ACL. Maximum 55 characters.

Add/Edit ACE

Configuration > VPN > WebVPN > ACLs > Add/Edit ACLs

An Access Control Entry permits or denies access to specific URLs and services. You can configure multiple ACEs for an ACL. ACLs apply ACEs in priority order, acting on the first match.

Fields

- Action—Permits or denies access to the specific networks, subnets, hosts, and web servers specified in the Filter group box.
- Filter—Specifies a URL or an IP address to which you want to apply the filter (permit or deny user access).
  - URL—Applies the filter to the specified URL.
  - Protocols (unlabeled)—Specifies the protocol part of the URL address.
  - ://x—Specifies the URL of the Web page to which to apply the filter.
  - TCP—Applies the filter to the specified IP address, subnet, and port.
  - IP Address—Specifies the IP address to which to apply the filter.
  - Netmask—Lists the standard subnet mask to apply to the address in the IP Address box.
  - Service—Identifies the service (such as https, kerberos, or any) to be matched. Displays a list of services from which you can select the service to display in the Service box.
  - Boolean operator (unlabeled)—Lists the boolean conditions (equal, not equal, greater than, less than, or range) to use in matching the service specified in the service box.
- Rule Flow Diagram—Graphically depicts the traffic flow using this filter. This area might be hidden.
- Options—Specifies the logging rules. The default is Default Syslog.
  - Logging—Choose enable if you want to enable a specific logging level.
  - Syslog Level—Grayed out until you select Enable for the Logging attribute. Lets you select the type of syslog messages you want the security appliance to display.
  - Log Interval—Lets you select the number of seconds between log messages.
  - Time Range—Lets you select the name of a predefined time-range parameter set.

– ...—Click to browse the configured time ranges or to add a new one.

Examples

Here are examples of WebVPN ACLs:

Action	Filter	Effect
Deny	url http://*.yahoo.com/	Denies access to all of Yahoo!
Deny	url cifs://fileserver/share/directory	Denies access to all files in the specified location.
Deny	url https://www.company.com/directory/file.html	Denies access to the specified file.
Permit	url https://www.company.com/directory	Permits access to the specified location
Deny	url http://*:8080/	Denies HTTPS access to anywhere via port 8080.
Deny	url http://10.10.10.10	Denies HTTP access to 10.10.10.10.
Permit	url any	Permits access to any URL. Usually used after an ACL that denies url access.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

APCF

Configuration > VPN > WebVPN > APCF

WebVPN includes an Application Profile Customization Framework option that lets the security appliance handle non-standard applications and web resources so they display correctly over a WebVPN connection. An APCF profile contains a script that specifies when (pre, post), where (header, body, request, response), and what data to transform for a particular application. The script is in XML and uses sed (stream editor) syntax to transform strings/text.

You can configure multiple APCF profiles on a security appliance to run in parallel. Within an APCF profile script, multiple APCF rules can apply. In this case, the security appliance processes the oldest rule first, based on configuration history, the next oldest rule next, and so forth.

You can store APCF profiles on the security appliance flash memory, or on an HTTP, HTTPS, FTP, or TFTP server. Use this panel to add, edit, and delete APCF packages, and to put them in priority order.

Fields

- APCF File Location—Displays information about the location of the APCF package. This can be on the security appliance flash memory, or on an HTTP, HTTPS, FTP, or TFTP server.
- Add/Edit—Click to add or edit a new or existing APCF profile.

- Delete—Click to remove an existing APCF profile. There is no confirmation or undo.
- Move Up/Move Down—Click to rearrange APCF profiles within a list. This determines the order in which the security appliance attempts to use APCF profiles.

Add/Edit APCF Profile

Configuration > VPN > WebVPN > APCF > Add/Edit APCF Profile

This panel lets you add or edit and APCF package, which includes identifying its location, which can be either on the security appliance flash memory, or on an HTTP, HTTPS, or TFTP server.

Fields

- Flash file—Check to locate an APCF file stored on the security appliance flash memory.
- Path—Displays the path to an APCF file stored on flash memory after you browse to locate it. You can also manually enter the path in this field.
- Browse Flash—Click to browse flash memory to locate the APCF file. A Browse Flash Dialog panel displays. Use the Folders and Files columns to locate the APCF file. Highlight the APCF file and click **OK.** The path to the file then displays in the Path field.

Note If you do not see the name of an APCF file that you recently downloaded, click the Refresh button.

- Upload —Click to upload an APCF file from a local computer to the security appliance flash file system. The Upload APCF package pane displays.
- URL—Check to use an APCF file stored on an HTTP, HTTPS or TFTP server.
- http/https/tftp (unlabeled)—Identify the server type.
- URL (unlabeled)—Enter the path to the HTTP, HTTPS, or TFTP server.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Upload APCF package

Configuration > VPN > WebVPN > APCF > Upload APCF Package

Fields

- Local File Path—Shows the path to the APCF file on your computer. Click **Browse Local** to automatically insert the path in this field, or enter the path.

- Browse Local—Click to locate and choose the APCF file on your computer that you want to transfer. The Select File Path dialog box displays the contents of the folder you last accessed on your local computer. Navigate to the APCF file, select it, and click **Open**. ASDM inserts the file path into the Local File Path field.
- Flash File System Path—Displays the path on the security appliance to upload the APCF file.
- Browse Flash—Click to identify the location on the security appliance to which you want to upload the APCF file. The Browse Flash dialog box displays the contents of flash memory.
- File Name—Located in the Browse Flash dialog box that opens when you click Browse Flash, this field displays the name of the APCF file you selected on your local computer. We recommend that you use this name to prevent confusion. Confirm that this file displays the correct filename, and click OK. The Browse Flash dialog box closes. ASDM inserts the destination file path in the Flash File System Path field.
- Upload File—Click when you have identified the location of the APCF file on your computer, and the location where you want to download it to the security appliance.
- A Status window appears and remains open for the duration of the file transfer. Following the transfer, an Information window displays the message, "File is uploaded to flash successfully." Click **OK**. The Upload Image dialog window removes the contents of the Local File Path and Flash File System Path fields, indicating you can upload another file. To do so, repeat these instructions. Otherwise, click the **Close** button.
- Close—Closes the Upload Image dialog window. Click this button after you upload the APCF file to flash memory or if you decide not to upload it. If you do upload it, the filename appears in the APCF File Location field of the APCF window. If you do not upload it, a Close Message dialog box prompts, "Are you sure you want to close the dialog without uploading the file?" Click **OK** if you do not want to upload the file. The Close Message and Upload Image dialog boxes close, revealing the APCF Add/Edit pane. Otherwise, click **Cancel** in the Close Message dialog box. The dialog box closes, revealing the Upload Image dialog box again, with the values in the fields intact. Click **Upload File**.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Auto Signon

Configuration > VPN > WebVPN > Auto Signon

The Auto Signon window or tab lets you configure or edit auto signon for WebVPN users. Auto signon is a simplified single signon method that you can use if you do not already have an SSO method deployed on your internal network. With auto signon configured for particular internal servers, the security appliance passes the login credentials that the WebVPN user used to login to the security appliance (username and password) to those particular internal servers. You configure the security appliance to

respond to a specific authentication method for a particular range of servers. The authentication methods you can configure the security appliance to respond to are NTLM authentication, HTTP Basic authentication, or both methods.

Auto signon is a straight-forward method for configuring SSO for particular internal servers. This section describes the procedure for setting up SSO with auto signon. If you already have SSO deployed using Computer Associates' SiteMinder SSO server and want to configure the security appliance to support this solution, see SSO Servers. If you use SSO with HTTP Forms protocol and want to configure the security appliance to support this method, see AAA Setup.

Fields

- IP Address—*Display only.* In conjunction with the following Mask, displays the IP address range of the servers to be authenticated to as configured with the Add/Edit Auto Signon dialog box. You can specify a server using either the server URI or the server IP address and mask.
- Mask—*Display only.* In conjunction with the preceding IP Address, displays the IP address range of the servers configured to support auto signon with the Add/Edit Auto Signon dialog box.
- URI—*Display only.* Displays a URI mask that identifies the servers configured with the Add/Edit Auto Signon dialog box.
- Authentication Type—*Display only.* Displays the type of authentication—basic HTTP, NTLM, or basic and NTLM—as configured with the Add/Edit Auto Signon dialog box.
- Add/Edit—Click to add or edit an auto signon instruction. An auto signon instruction defines a range of internal servers using the auto signon feature and the particular authentication method.
- Delete—Click to delete an auto signon instruction selected in the Auto Signon table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Auto Signon Entry

You can get to this panel through various paths.

The Add/Edit Auto Signon Entry dialog box lets you add or edit a new auto signon instruction. An auto signon instruction defines a range of internal servers using the auto signon feature and the particular authentication method.

Fields

- IP Block—Click this button to specify a range of internal servers using an IP address and mask.
  - IP Address—Enter the IP address of the first server in the range for which you are configuring auto-signon.
  - Mask—In the subnet mask menu, click the subnet mask that defines the server address range of the servers supporting auto signon.

- URI—Click this button to specify a server supporting auto signon by URI, then enter the URI in the field next to this button.
- Authentication Type—The authentication method assigned to the servers. For the specified range of servers, the security appliance can be configured to respond to HTTP Basic authentication requests, NTLM authentication requests, or requests using either method.
  - Basic—Click this button to assign basic HTTP authentication.
  - NTLM—Click this button use NTLMv1 authentication.
  - Basic and NTLM—Click this button use either HTTP Basic or NTLMv1 authentication.

Note If you configure one method for a range of servers (e.g., HTTP Basic) and one of those servers attempts to authenticate with a different method (e.g., NTLM), the security appliance does not pass the users login credentials to that server,

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

CSD Setup

Configuration > VPN > WebVPN > CSD Setup

This window lets you view the version and state of the CSD distribution package, and install, upgrade, enable, and disable CSD.

Fields

CSD Setup

- Secure Desktop Image—Displays the CSD distribution package loaded into the running configuration. This field should display the filename in the format securedesktop_asa_<n>_<n>*.pkg. Use the Browse Flash button to insert or modify the value in this field. You can also use this field to view the version of CSD. (The Configuration CSD Secure Desktop Manager also displays the CSD version.)
- Enable Secure Desktop—Check and click **Apply** to do the following:
  a. Make sure the file is a valid CSD distribution package.
  b. Create an “sdesktop” folder on disk0 if one is not already present.
  c. Insert a data.xml (CSD configuration) file into the sdesktop folder if one is not already present.
  d. Load the data.xml file into the running configuration.

Note If you transfer or replace the data.xml file, disable and then enable CSD to load the file.

e. Enable CSD.

- Browse Flash—Click to view the contents of the flash device and choose or type the filename of the CSD distribution package to install into the running configuration. You can use this button to install, upgrade or downgrade CSD. Click **Apply** to save the CSD setup.

Note If you click the **Browse Flash** button to upgrade or downgrade the CSD distribution package, select the package to install, and click **OK**, the Uninstall CSD dialog window asks you if you want to delete the CSD distribution currently in the running configuration from the flash device. Click **Yes** if you want to save space on the flash device, or click **No** to reserve the option to revert to this version of CSD.

- Upload—Lets you transfer a copy of a CSD distribution package from your local computer to the flash device. To prepare to install or upgrade CSD, use your Internet browser to download a securedesktop_asa_<*n*>_<*n*>*.pkg file from http://www.cisco.com/cgi-bin/tablebuild.pl/securedesktop to any location on your PC. Then use this button to transfer a copy from you local computer to the flash device. Finally, click **Browse Flash** to install it into the running configuration.
- Uninstall—Lets you remove the CSD image and configuration file (sdesktop/data.xml) from the running configuration. If you click this button, the Uninstall CSD dialog window asks if you want to delete the CSD image that was named in the "Secure Desktop Image field" and all CSD data files (including the entire CSD configuration) from the flash device. Click **Yes** if you want to remove these files from both the running configuration and the flash device, or click **No** to remove them from the running configuration, but retain them on the flash device.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Upload Image

Configuration > VPN > WebVPN > CSD Setup > Upload

This dialog window lets you transfer a copy of a CSD distribution package from your local computer to the flash device on the security appliance. Use this window to install or upgrade CSD.

Note Before using this window, use your Internet browser to download a securedesktop_asa_<*n*>_<*n*>*.pkg file from http://www.cisco.com/cgi-bin/tablebuild.pl/securedesktop to any location on your local computer.

Fields

Use the fields and options in this window as follows:

- Local File Path—Specifies the path to the securedesktop_asa_<*n*>_<*n*>*.pkg file on your local computer. Click **Browse Local** to automatically insert the path in this field, or enter the path. For example:

  D:\Documents and Settings*Windows_user_name*.AMER\My Documents\My Downloads\securedesktop_asa_3_1_1_16.pkg

- Browse Local—Click to select the path of the securedesktop_asa_<*n*>_<*n*>*.pkg file to be transferred. The Selected File Path dialog box displays the contents of the folder you last accessed on your local computer. Navigate to the securedesktop_asa_<*n*>_<*n*>*.pkg file, select it, and click **Open**.

  ASDM inserts the file path into the Local File Path field.

- Flash File System Path—Specifies the destination path on the flash device of the security appliance and the name of the destination file. Click **Browse Flash** to automatically insert the path into this field, or enter the path. For example,

  disk0:/securedesktop_asa_3_1_1_16.pkg

- Browse Flash—Click to select the target directory for the file. The Browse Flash dialog box displays the contents of the flash card.

- File Name—Located in the Browse Flash dialog box that opens if you click **Browse Flash**, this field displays the name of the CSD distribution package you selected on your local computer. We recommend that you use this name to prevent confusion. Confirm that this field displays the same name of the local file you selected and click **OK**. The Browse Flash dialog box closes. ASDM inserts the destination file path into the Flash File System Path field.

- Upload File—Uploads the securedesktop_asa_<*n*>_<*n*>*.pkg file from your local computer to the flash device. A Status window appears and remains open for the duration of the file transfer. Following the transfer, an Information window displays the message, “File is uploaded to flash successfully.” Click **OK**. The Upload Image dialog window removes the contents of the Local File Path and Flash File System Path fields, indicating you can upload another file. To do so, repeat these instructions. Otherwise, click the **Close** button.

- Close—Closes the Upload Image dialog window. Click this button after you upload the CSD distribution package to the flash device or if you decide not to upload it. If you uploaded it, the filename appears in the Secure Desktop Image field of the CSD Setup window. If you did not upload it, a Close Message dialog box prompts, “Are you sure you want to close the dialog without uploading the file?” Click **OK** if you do not want to upload the file. The Close Message and Upload Image dialog boxes close, revealing the CSD Setup pane. Otherwise, click **Cancel** in the Close Message dialog box. The dialog box closes, revealing the Upload Image dialog box again, with the values in the fields intact. Click **Upload File**.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Cache

Configuration > VPN > WebVPN > Cache

Caching enhances WebVPN performance. It stores frequently reused objects in the system cache, which reduces the need to perform repeated rewriting and compressing of content. It reduces traffic between WebVPN and both the remote servers and end-user browsers, with the result that many applications run much more efficiently.

Fields

- Enable cache—Check to enable caching.
- Parameters—Lets you define the terms for caching.
  - Enable caching of compressed content—Check to cache compressed content. When you disable this parameter, the security appliance stores objects before it compresses them.
  - Maximum Object Size—Enter the maximum size in KB of a document that the security appliance can cache. The security appliance measures the original content length of the object, not rewritten or compressed content. The range is 0 to 10,000 KB; the default is 1000 KB
  - Minimum Object Size—Enter the minimum size in KB of a document that the security appliance can cache. The security appliance measures the original content length of the object, not rewritten or compressed content. The range is 0 to 10,000 KB; the default is 0 KB.

Note The Maximum Object Size must be greater than the Minimum Object Size.

  - LM Factor—Enter an integer between 1 and 100; the default is 20.

    The LM factor sets the policy for caching objects which have only the last-modified timestamp. This revalidates objects that have no server-set change values. The security appliance estimates the length of time since the object has changed, also called the expiration time. The estimated expiration time equals the time elapsed since the last change multiplied by the LM factor. Setting the LM factor to 0 forces immediate revalidation, while setting it to 100 results in the longest allowable time until revalidation.
  - Expiration Time—Enter an integer between 0 and 900 to set the number of minutes to cache objects without revalidating them. The default is one minute.

    The expiration time sets the amount of time to for the security appliance to cache objects that have neither a last-modified time stamp nor an explicit server-set expiry time.
- Restore Cache Default—Click to restore default values for all cache parameters.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Content Rewrite

Configuration > VPN > WebVPN > Content Rewrite

The Content Rewrite panel lists all applications for which content rewrite is enabled or disabled.

WebVPN processes application traffic through a content transformation/rewriting engine that includes advanced elements such as JavaScript, VBScript, Java, and multi-byte characters to proxy HTTP traffic which may have different semantics and access control rules depending on whether the user is using an application within or independently of an SSL VPN device.

You might not want some applications and web resources, for example, public websites, to go through the security appliance. The security appliance therefore lets you create rewrite rules that let users browse certain sites and applications without going through the security appliance. This is similar to split-tunneling in an IPSec VPN connection.

You can create multiple rewrite rules. The rule number is important because the security appliance searches rewrite rules by order number, starting with the lowest, and applies the first rule that matches.

Fields

- Content Rewrite
  - Rule Number—Displays an integer that indicates the position of the rule in the list.
  - Rule Name—Provides the name of the application for which the rule applies.
  - Rewrite Enabled—Displays content rewrite as enabled or disabled.
  - Resource Mask—Displays the resource mask.
- Add/Edit—Click to add a rewrite entry or edit a selected rewrite entry.
- Delete—Click to delete a selected rewrite entry.

.Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Content Rewrite Rule

Configuration > VPN > WebVPN > Content Rewrite > Add/Edit Content Rewrite Rule

- Enable content rewrite—Check to enable content rewrite for this rewrite rule.
- Rule Number—(Optional) Enter a number for this rule. This number specifies the position of the rule in the list. Rules without a number are at the end of the list. The range is 1 to 65534.
- Rule Name—(Optional) Provide an alphanumeric string that describes the rule, maximum 128 characters.
- Resource Mask—Enter the resource mask. This is a word, length up to 300 characters.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Java Trustpoint

Configuration > VPN > WebVPN > Java Trustpoint

Java objects which have been transformed by WebVPN can subsequently be signed using a PKCS12 digital certificate associated with a trustpoint. In the Java Trustpoint pane, you can configure the WebVPN Java object signing facility to use a PKCS12 certificate and keying material from a specified trustpoint location. To import a trustpoint, see Configuration > Properties > Certificate > Trustpoint > Import.

Fields

- Java Trustpoint—Choose the configured trustpoint that you want to employ in Java object signing.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Encoding

Configuration > VPN > WebVPN > Encoding

This window lets you specify the character encoding for WebVPN portal pages to remote clients.

Character encoding, also called “character coding” and “a character set,” is the pairing of raw data (such as 0’s and 1’s) with characters to represent the data. The language determines the character encoding method to use. Some languages use the same method, while others do not. Usually, the geographic region determines the default encoding method used by the browser, but the remote user can change this. The browser can also detect the encoding specified on the page, and render the document accordingly.

The encoding attribute lets you specify the value of the character-encoding method into the WebVPN portal page to ensure that the browser renders it properly, regardless of the region in which the user is using the browser, or any changes made to the browser.

By default, the security appliance applies the "Global WebVPN Encoding Type" to pages from Common Internet File System servers. The mapping of CIFS servers to their appropriate character encoding, globally with the "Global WebVPN Encoding Type" attribute, and individually with the file-encoding exceptions displayed in the table, provides for the accurate handling and display of CIFS pages when the proper rendering of filenames or directory paths, as well as pages, are an issue.

Fields

- Global WebVPN Encoding Type —This attribute determines the character encoding that all WebVPN portal pages inherit except for those from the CIFS servers listed in the table. You can type the string, or select one from the drop-down list, which contains only the most common values, as follows:
  - big5
  - gb2312
  - ibm-850
  - iso-8859-1
  - shift_jis

  Note If you are using Japanese Shift_jis Character encoding, click **Do not specify** in the Font Family area of the associated Select Page Font pane to remove the font family.

  - unicode
  - windows-1252
  - none

    If you choose **none** or specify a value that the browser on the WebVPN client does not support, it uses its own default encoding.

  You can type a string consisting of up to 40 characters, and equal to one of the valid character sets identified in http://www.iana.org/assignments/character-sets. You can use either the name or the alias of a character set listed on that page. The string is case-insensitive. The command interpreter converts upper-case to lower-case when you save the security appliance configuration.
- CIFS Server—Name or IP address of each CIFS server for which the encoding requirement differs from the "Global WebVPN Encoding Type" attribute setting.

  A difference in the encoding of the CIFS server filename and directory indicates that you might need to add an entry for the server to ensure the encoding is correct.
- Encoding Type—Displays the character encoding override for the associated CIFS server.
- Add—Click once for each CIFS server for which you want to override the "Global WebVPN Encoding Type" setting.
- Edit—Select a CIFS server in the table and click this button to change its character encoding.
- Delete—Select a CIFS server in the table and click this button to delete the associated entry from the table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add\Edit Encoding

Configuration > VPN > WebVPN > Encoding > Add\Edit

This dialog window lets you maintain exceptions to the "Global WebVPN Encoding Type" attribute setting in the Configuration > VPN > WebVPN > Encoding window. That window contains the Add and Edit buttons that open this dialog box.

Fields

- CIFS Server—Enter the name or IP address of a CIFS server for which the encoding requirement differs from the "Global WebVPN Encoding Type" attribute setting. The security appliance retains the case you specify, although it ignores the case when matching the name to a server.
- Encoding Type —Choose the character encoding that the CIFS server should provide for WebVPN portal pages. You can type the string, or select one from the drop-down list, which contains only the most common values, as follows:
  - big5
  - gb2312
  - ibm-850
  - iso-8859-1
  - shift_jis

  Note If you are using Japanese Shift_jis Character encoding, click **Do not specify** in the Font Family area of the associated Select Page Font pane to remove the font family.

  - unicode
  - windows-1252
  - none

    If you choose **none** or specify a value that the browser on the WebVPN client does not support, it uses its own default encoding.

  You can type a string consisting of up to 40 characters, and equal to one of the valid character sets identified in http://www.iana.org/assignments/character-sets. You can use either the name or the alias of a character set listed on that page. The string is case-insensitive. The command interpreter converts upper-case to lower-case when you save the security appliance configuration.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Port Forwarding

Configuration > VPN > WebVPN > Port Forwarding

Port forwarding lets users access TCP-based applications over a WebVPN connection. Such applications include the following:

Lotus Notes	Secure FTP (FTP over SSH)
Outlook Express	SSH
Outlook	Telnet
Perforce	Windows Terminal Service
Sametime	XDDTS

Other TCP-based applications may also work, but we have not tested them. Protocols that use UDP do not work.

Note Port forwarding supports only those TCP applications that use static TCP ports. Applications that use dynamic ports or multiple TCP ports are not supported. For example, SecureFTP, which uses port 22, works over WebVPN port forwarding, but standard FTP, which uses ports 20 and 21, does not.

Port forwarding does not support connections to personal digital assistants.

Port Forwarding and JRE

Because port forwarding requires downloading the Java applet and configuring the local client, and because doing so requires administrator permissions on the local system, it is unlikely that users will be able to use applications when they connect from public remote systems.

Caution Make sure Sun Microsystems Java™ Runtime Environment (JRE) 1.5.x is installed on the remote computers to support port forwarding (application access) and digital certificates. If JRE 1.4.x is running and the user authenticates with a digital certificate, the application fails to start because JRE cannot access the web browser's certificate store.

The Java applet displays in its own window on the end user HTML interface. It shows the contents of the list of forwarded ports available to the user, as well as which ports are active, and amount of traffic in bytes sent and received.

Port Forwarding and User Authentication Via Digital Certificates Incompatibility

Neither port forwarding nor the ASDM JAVA applet work with user authentication using digital certificates. JAVA does not have the ability to access the web browser keystore. Therefore JAVA cannot use certificates that the browser uses to authenticate users, and the application cannot start.

Fields

- Configure port forwarding lists for application access over WebVPN group—To configure application access, create one or more named lists of applications, and then assign a list, by name, to a user (Configuration > Properties > Device Administration > User Accounts > Add/Edit User Account / WebVPN tab) or a group policy (Configuration > VPN > General > Add/Edit Group Policy > WebVPN tab). You can associate a user or group policy with one list only.
  - List Name—Displays the names of application lists configured for WebVPN.
  - Local TCP Port—Displays the local port that listens for traffic for the application.
  - Remote Server—Displays the IP address or DNS name of the remote server.
  - Remote TCP Port—Displays the remote port that listens for traffic for the application.
  - Description—Displays text that describes the TCP application.
- Add/Edit—Click to add or modify a port forwarding list.
- Delete—Click to remove an existing port forwarding list. There is no confirmation or undo.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

For More Information

WebVPN End User Set-up

Add/Edit Port Forwarding List

Configuration > VPN > WebVPN > Port Forwarding > Add/Edit Port Forwarding List

The Add/Edit Port Forwarding List panels let you add or edit a named list of TCP applications to associate with users or group policies for access over WebVPN connections.

Fields

- List Name—Enter an alpha-numeric name for the list. Maximum 64 characters.
  - Local TCP Port—Displays the local port that listens for traffic for the application.
  - Remote Server—Displays the IP address or DNS name of the remote server.
  - Remote TCP Port—Displays the remote port that listens for traffic for the application.
  - Description—Displays text that describes the TCP application.

- Add/Edit—Click to add or modify a port forwarding list.
- Delete—Click to remove an existing port forwarding list. There is no confirmation or undo.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Port Forwarding Entry

Configuration > VPN > WebVPN > Port Forwarding > Add/Edit Port Forwarding List > Add/Edit Port Forwarding Entry

The Add/Edit Port Forwarding Entry panels let you configure specific applications for a named port forwarding list.

Fields

- Local TCP Port—Type a port number for the application to use. You can use a local port number only once for a listname. To avoid conflicts with local TCP services, use port numbers in the range 1024 to 65535.
- Remote Server—Type either the DNS name or IP address of the remote server. We recommend using hostnames so that you do not have to configure the client applications for specific IP addresses.
- Remote TCP Port—Type the well-know port number for the application.
- Description—Type a description of the application. Maximum 64 characters.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Proxies

Configuration > VPN > WebVPN > Proxies

The security appliance can terminate HTTPS connections and forward HTTP/HTTPS requests to HTTP and HTTPS proxy servers. These servers act as an intermediary between users and the Internet. Requiring all Internet access via a server you control provides another opportunity for filtering to assure secure Internet access and administrative control.

Be aware that HTTP/HTTPS proxy does not support connections to personal digital assistants.

Fields

- HTTP—Lets you define an HTTP proxy server.
  - IP Address—Enter the IP address of the HTTP proxy server.
  - Port—Enter the port that listens for HTTP requests. The default port is 80.
- HTTPS—Lets you define an HTTPS proxy server.
  - IP Address—Enter the IP address of the HTTPS proxy server.
  - Port—Enter the port that listens for HTTPS requests. The default port is 443.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Proxy Bypass

Configuration > VPN > WebVPN > Proxy Bypass

You can configure the security appliance to use proxy bypass when applications and web resources work better with the special content rewriting this feature provides. Proxy bypass is an alternative method of content rewriting that makes minimal changes to the original content. It is often useful with custom web applications.

You can configure multiple proxy bypass entries. The order in which you configure them is unimportant. The interface and path mask or interface and port uniquely identify a proxy bypass rule.

If you configure proxy bypass using ports rather than path masks, depending on your network configuration, you might need to change your firewall configuration to allow these ports access to the security appliance. Use path masks to avoid this restriction. Be aware, however, that path masks can change, so you might need to use multiple pathmask statements to exhaust the possibilities.

A path is the text in a URL that follows the domain name. For example, in the URL www.mycompany.com/hrbenefits, *hrbenefits* is the path. Similarly, for the URL www.mycompany.com/hrinsurance, *hrinsurance* is the path. If you want to use proxy bypass for all hr sites, you can avoid using the command multiple times by using the * wildcard as follows: /hr*.

Fields

- Interface—Displays the VLAN configured for proxy bypass.
- Port—Displays the port configured for proxy bypass.
- Path Mask—Displays the URI path to match for proxy bypass.
- URL—Displays the target URLs.
- Rewrite—Displays the rewrite options. These are a combination of XML, link, or none.
- Add/Edit—Click to add a proxy bypass entry or edit a selected entry.

- Delete—Click to delete a proxy bypass entry.

.Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Proxy Bypass Rule

Configuration > VPN > WebVPN > Proxy Bypass > Add/Edit Proxy Bypass Rule

This panel lets you set rules for when the security appliance performs little or no content rewriting.

Fields

- Interface Name—Select the VLAN for proxy bypass.
- Bypass Condition—Specify either a port or a URI for proxy bypass.
  - Port—Select to use a port for proxy bypass. Valid port numbers are 20000-21000.
  - Port (unlabeled)—Enter a high-numbered port for the security appliance to reserve for proxy bypass.
  - Path Mask—Select to use a URL for proxy bypass.
  - Path Mask—Enter a URL for proxy bypass. It can contain a regular expression.
- URL—Define target URLs for proxy bypass.
  - Protocol—Select either http or https as the protocol.
  - URL (unlabeled)—Enter a URL to which you want to apply proxy bypass.
- Content to Rewrite—Specifies the content to rewrite. The choices are none or a combination of XML, links, and cookies.
  - XML—Check to rewrite XML content.
  - Hostname—Check to rewrite links.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

SSL VPN Client

Configuration > VPN > WebVPN > SSL VPN Client

This window lets you enable the security appliance to download SVC image files to remote computers. The SVC Image Files pane displays files existing in flash memory identified as SVC images. The order of the files in this pane indicates the order in which they are downloaded to the remote computer.

SVC is a VPN tunneling technology that gives remote users the benefits of an IPSec VPN client without the need for network administrators to install and configure IPSec VPN clients on remote computers. The SVC uses the SSL encryption that is already present on the remote computer as well as the WebVPN login and authentication of the security appliance.

To establish an SVC session, the remote user enters the IP address of a WebVPN interface of the security appliance in the browser, and the browser connects to that interface and displays the WebVPN login screen. If the user satisfies the login and authentication, and the security appliance identifies the user as *requiring* the SVC, the security appliance downloads the SVC to the remote computer. If the security appliance identifies the user as having the *option* to use the SVC, the security appliance downloads the SVC to the remote computer while presenting a link on the user screen to skip the SVC installation.

After downloading, the SVC installs and configures itself, and then the SVC either remains or uninstalls itself (depending on the configuration) from the remote computer when the connection terminates.

The security appliance might have several unique SVC images residing in cache memory for different remote computer operating systems. When the user attempts to connect, the security appliance can consecutively download portions of these images to the remote computer until the image and operating system match, at which point it downloads the entire SVC. You can order the SVC images to minimize connection setup time, with the first image downloaded representing the most commonly-encountered remote computer operating system.

Fields

- Enable—Enables the security appliance to download SVC image files to remote computers.
- Add—Displays the Add SSL VPN Client Image window, where you can specify a file in flash memory as an SVC image file, or where you can browse flash memory for a file to specify as an SVC image. You can also upload a file from a local computer to the flash memory.
- Replace—Displays the Replace SSL VPN Client Image window, where you can specify a file in flash memory as an SVC image to replace an SVC image hightlighted in the SVC Image Files table. You can also upload a file from a local computer to the flash memory.
- Delete—Deletes an SVC image that you hightlight in the SVC Image Files pane.
- Move Up and Move Down—changes the order in which the security appliance downloads the SVC images to the remote computer. It downloads the SVC image at the top of the SVC Image Files pane first. Therefore, you should move the SVC image used by the most commonly-encountered operating system to the top of the pane.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

For More Information

WebVPN End User Set-up

Add SSL VPN Client Image

Configuration > VPN > WebVPN > SSL VPN Client> Add SSL VPN Client Image

In this window, you can specify a filename for a file on the security appliance flash memory that you want to identify as an SSL VPN Client (SVC) image. You can also browse the flash memory for a file to identify, or you can upload a file from a local computer to the flash memory.

Fields

- Flash SVC Image—Specify the filename of the file in flash memory that you want to identify as an SSL VPN Client (SVC) image.
- Browse Flash—Displays the Browse Flash Dialog window where you can view all the files on flash memory.
- Upload—Displays the Upload Image window where you can upload a file from a local PC that you want to identify as an SVC image.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add SSL VPN Client Browse Flash Dialog

Configuration > VPN > WebVPN > SSL VPN Client> Add SSL VPN Client Image > Browse Flash Dialog

In this window, you can browse the flash memory of the security appliance for a file that you want to identify as an SSL VPN Client (SVC) image. You can also browse the flash memory for a file to identify, or you can upload a file from a local computer to the flash memory.

Fields

- Flash SVC Image—Identifies the filename of the file in flash memory that you want to identify as an SSL VPN Client (SVC) image.
- Browse Flash—Displays the Browse Flash Dialog window where you can view all the files on flash memory.
- Upload—Displays the Upload Image window where you can upload a file from a local PC that you want to identify as an SVC image.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add SSL VPN Client Upload Flash Dialog

Configuration > VPN > WebVPN > SSL VPN Client> Add SSL VPN Client Image > Browse Flash Dialog

In this window, you can specify the path of a file on the local computer or in flash memory of the security appliance that you want to identify as an SSL VPN Client (SVC) image. You can also browse the local computer or the flash memory of the security appliance for a file to identify.

Fields

- Local File Path—Identifies the filename of the file in on the local computer that you want to identify as an SSL VPN Client (SVC) image.
- Browse Local—Displays the Select File Path window where you can view all the files on local computer and where you can select a file to identify as an SVC image.
- Flash File System Path—Identifies the filename of the file in the flash memory of the security appliance that you want to identify as an SSL VPN Client (SVC) image.
- Browse Flash—Displays the Browse Flash Dialog window where you can view all the files on flash memory of the security appliance and where you can select a file to identify as an SVC image.

Replace SSL VPN Client Image

Configuration > VPN > WebVPN > SSL VPN Client> Replace SSL VPN Client Image

In this window, you can specify a filename for a file on the security appliance flash memory that you want to identify as an SSL VPN Client (SVC) image to replace a file previously identified as an SVC image. You can also browse the flash memory for a file to identify, or you can upload a file from a local computer to the flash memory.

Fields

- Flash SVC Image—Specify the filename of the file in flash memory that you want to identify as an SSL VPN Client (SVC) image.
- Browse Flash—Displays the Browse Flash Dialog window where you can view all the files on flash memory.
- Upload—Displays the Upload Image window where you can upload a file from a local PC that you want to identify as an SVC image.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Replace SSL VPN Client Upload Flash Dialog

Configuration > VPN > WebVPN > SSL VPN Client > Replace SSL VPN Client Image > Browse > Flash Dialog

In this window, you can specify the path of a file on the local computer or in flash memory of the security appliance that you want to identify as an SSL VPN Client (SVC) image. You can also browse the local computer or the flash memory of the security appliance for a file to identify.

Fields

- Local File Path—Identifies the filename of the file in on the local computer that you want to identify as an SSL VPN Client (SVC) image.
- Browse Local—Displays the Select File Path window where you can view all the files on local computer and where you can select a file to identify as an SVC image.
- Flash File System Path—Identifies the filename of the file in the flash memory of the security appliance that you want to identify as an SSL VPN Client (SVC) image.
- Browse Flash—Displays the Browse Flash Dialog window where you can view all the files on flash memory of the security appliance and where you can select a file to identify as an SVC image.

SSO Servers

Configuration > VPN > WebVPN > SSO Servers

The SSO Server window lets you configure or delete single sign-on (SSO) for WebVPN users using Computer Associates' SiteMinder SSO server. SSO support, available only for WebVPN, lets users access different secure services on different servers without entering a username and password more than once.

You can choose from three methods when configuring SSO: Auto Signon using basic HTTP and/or NTLMv1 authentication, HTTP Form protocol, or Computer Associates eTrust SiteMinder (formerly Netegrity SiteMinder). This section describes the procedure for setting up SSO with SiteMinder.

- To configure SSO with basic HTTP or NTLM authentication, see Auto Signon.
- To configure SSO with the HTTP Form protocol, see AAA Setup.

The SSO mechanism either starts as part of the AAA process (HTTP Forms) or just after successful user authentication to a AAA server (SiteMinder). In both cases, the WebVPN server running on the security appliance acts as a proxy for the user to the authenticating server. When a user logs in, the WebVPN server sends an SSO authentication request, including username and password, to the authenticating server using HTTPS. If the server approves the authentication request, it returns an SSO authentication cookie to the WebVPN server. This cookie is kept on the security appliance on behalf of the user and used to authenticate the user to secure websites within the domain protected by the SSO server.

SSO authentication with SiteMinder is separate from AAA and occurs after the AAA process completes. To set up SSO for a user or group, you must first configure a AAA server (RADIUS, LDAP and so forth). After a user authenticates to the AAA server, the WebVPN server uses HTTPS to send an authentication request to the SiteMinder SSO server.

Besides configuring the security appliance, you must also configure your CA SiteMinder Policy Server with the Cisco authentication scheme. See Adding the Cisco Authentication Scheme to SiteMinder.

Fields

- Server Name—*Display only.* Displays the names of configured SSO Servers. The minimum number of characters is 4, and the maximum is 31.
- Authentication Type—*Display only.* Displays the type of SSO server. The security appliance currently supports the SiteMinder type.
- URL—*Display only.* Displays the SSO server URL to which the security appliance makes SSO authentication requests.
- Secret Key—*Display only.* Displays the secret key used to encrypt authentication communications with the SSO server. The key can be comprised of any regular or shifted alphanumeric character. There is no minimum or maximum number of characters.
- Maximum Retries—*Display only.* Displays the number of times the security appliance retries a failed SSO authentication attempt. The range is 1 to 5 retries, and the default number of retries is 3.
- Request Timeout (seconds)—*Display only.* Displays the number of seconds before a failed SSO authentication attempt times out. The range is 1 to 30 seconds, and the default number of seconds is 5.
- Add/Edit—Opens the Add/Edit SSO Server dialog box.
- Delete—Deletes the selected SSO server.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Adding the Cisco Authentication Scheme to SiteMinder

Besides configuring the security appliance for SSO with SiteMinder, you must also configure your CA SiteMinder Policy Server with the Cisco authentication scheme, provided as a Java plug-in.

Note

- Configuring the SiteMinder Policy Server requires experience with SiteMinder.
- This section presents general tasks, not a complete procedure.
- Refer to the CA SiteMinder documentation for the complete procedure for adding a custom authentication scheme.

To configure the Cisco authentication scheme on your SiteMinder Policy Server, perform the following tasks:

Step 1 With the Siteminder Administration utility, create a custom authentication scheme being sure to use the following specific arguments:

- In the Library field, enter **smjavaapi**.
- In the Secret field, enter the same secret configured in the Secret Key field of the Add SSO Server dialog to follow.
- In the Parameter field, enter **CiscoAuthAPI**.

Step 2 Using your Cisco.com login, download the file **cisco_vpn_auth.jar** from http://www.cisco.com/cgi-bin/tablebuild.pl/asa and copy it to the default library directory for the SiteMinder server.

Add/Edit SSO Server

Configuration > VPN > WebVPN > SSO Servers > Add/Edit SSO Server

Note This SSO method uses CA SiteMinder. You can also set up SSO using the HTTP Form protocol, or Basic HTML and NTLM authentication. To use the HTTP Form protocol, see AAA Setup. To set use basic HTML or NTLM authentication, use the **auto-signon** command at the command line interface.

Fields

- Server Name—If adding a server, enter the name of the new SSO server. If editing a server, this field is display only; it displays the name of the selected SSO server.
- Authentication Type—*Display only.* Displays the type of SSO server. The type currently supported by the security appliance is SiteMinder.
- URL—Enter the SSO server URL to which the security appliance makes SSO authentication requests.
- Secret Key—Enter a secret key used to encrypt authentication requests to the SSO server. Key characters can be any regular or shifted alphanumeric characters. There is no minimum or maximum number of characters. The secret key is similar to a password: you create it, save it, and configure it. It is configured on both the security appliance and the SiteMinder Policy Server using the Cisco Java plug-in authentication scheme.
- Maximum Retries—Enter the number of times the security appliance retries a failed SSO authentication attempt before the authentication times-out. The range is from 1 to 5 retries inclusive, and the default is 3 retries.
- Request Timeout—Enter the number of seconds before a failed SSO authentication attempt times out. The range is from1 to 30 seconds inclusive, and the default is 5 seconds.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Servers and URLs

Configuration > VPN > WebVPN > Servers and URLs

The Servers and URLs lets you view, add, and populate lists servers and URLs for access over WebVPN.

Note File access requires that you configure a NetBIOS server (Configuration > VPN > General > Tunnel Group > Add/Edit Tunnel Group > WebVPN > NetBIOS Servers).

Fields

Configure lists of servers and URLs for access over WebVPN—To configure file and URL access, create one or more named lists of file servers and URLs, and then assign the listname to individual users (Configuration > Properties > Device Administration > User Accounts > Add/Edit User Account / WebVPN tab > Other) or a group policy (Configuration > VPN > General > Add/Edit Group Policy > WebVPN tab > Other). You can associate a user or group policy with only one list.

- List Name—Names of server and URL lists configured for WebVPN.
- URL Display Name—Names end users see for the individual servers and URLs in the list.
- URL—URLs or paths to servers in the list.
- Add—Click to add a list of servers and URLs.
- Edit—Select a list in the Servers and URLs box and click this button to modify it.
- Delete—Select a list in the Servers and URLs box and click this button to remove it. ASDM removes the list without confirming the request.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

WebVPN Access

Configuration > VPN > WebVPN > WebVPN Access

The WebVPN Access panel lets you accomplish the following tasks:

- Enable or disable security appliance interfaces for WebVPN sessions
- Choose a port for WebVPN connections
- Set a global timeout value for WebVPN sessions
- Set a maximum number of simultaneous WebVPN sessions
- Configure the amount of security appliance memory that WebVPN can use.

To configure WebVPN services for individual users, the best practice is to use the **Configuration > VPN > General > Group Policy >Add/Edit >WebVPN** panel. Then use the **Configuration > Properties >Device Administration >User Accounts > VPN Policy** panel to assign the group policy to a user.

Fields

- Configure access parameters for WebVPN—Lets you enable or disable WebVPN connections on configured security appliance interfaces.
  - Interface—Displays names of all configured interfaces.
  - WebVPN Enabled—Displays current status for WebVPN on the interface.

    A green check next to Yes indicates that WebVPN is enabled.

    A red circle next to No indicates that WebVPN is disabled.
  - Enable/Disable—Click to enable or disable WebVPN on the highlighted interface.
- Port Number—Enter the port number that you want to use for WebVPN sessions. The default port is 443, for HTTPS traffic; the range is 1 through 65535. If you change the port number, All current WebVPN connections terminate, and current users must reconnect. You also lose connectivity to ASDM, and a prompt displays, inviting you to reconnect.
- Default Idle Timeout—Enter the amount of time, in seconds, that a WebVPN session can be idle before the security appliance terminates it. This value applies only if the Idle Timeout value in the group policy for the user is set to zero (0), which means there is no timeout value; otherwise the group policy Idle Timeout value takes precedence over the timeout you configure here. The minimum value you can enter is 1 minute. The default is 30 minutes (1800 seconds). Maximum is 24 hours (86400 seconds).

  We recommend that you set this attribute to a short time period. This is because a browser set to disable cookies (or one that prompts for cookies and then denies them) can result in a user not connecting but nevertheless appearing in the sessions database. If the Simultaneous Logins attribute for the group policy is set to one, the user cannot log back in because the database indicates that the maximum number of connections already exists. Setting a low idle timeout removes such phantom sessions quickly, and lets a user log in again.
- Max. Sessions Limit—Enter the maximum number of WebVPN sessions you want to allow. Be aware that the different ASA models support WebVPN sessions as follows: ASA 5510 supports a maximum of 150; ASA 5520 maximum is 750; ASA 5540 maximum is 2500.
- WebVPN Memory Size—Enter the percent of total memory or the amount of memory in kilobytes that you want to allocate to WebVPN processes. The default is 50% of memory. Be aware that the different ASA models have different total amounts of memory as follows: ASA 5510—256 MB; ASA5520 —512 MB: ASA 5540—1GB. When you change the memory size, the new setting takes effect only after the system reboots.
- WebVPN Memory (unlabeled)—Choose to allocate memory for WebVPN either as a percentage of total memory or as an amount of memory in kilobytes.

- Enable Tunnel Group Drop-down List on WebVPN Login— Check to include a drop-down list of configured tunnel groups on the WebVPN end-user interface. Users select a tunnel group from this list when they log on. This field is checked by default. If you uncheck it, the user cannot select a tunnel group at logon.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

For More Information

WebVPN End User Set-up

Webpage Customization

Configuration > VPN > WebVPN > Customization

Webpage Customization lets you customize the appearance of the WebVPN page that appears to WebVPN users when they connect to the security appliance. You can also customize pages that display to WebVPN users after the security appliance authenticates them, including the WebVPN Home page and the Application Access page.

Fields

- Customization Objects table—Displays the default WebVPN customization object (DfltCustomization), and any customization objects you add.
- Add—Adds a new customization object and displays the Add Customization Object dialog box where you can further customize.
- Edit—For the highlighted object in the Customization Object table, the Edit button displays the Edit Customization Object dialog box, where you can further customize.
- Delete—Deletes the highlighted object in the Customization Object table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Select Font

You can get to this dialog box from various paths.

The Select Font dialog box lets you specify a font family, style, weight, and size.

Fields

- Font Family—Lets you customize the font family.
  - Do not specify—Specifies the default.
  - Use selection—Enables a list of font families that you can select from.
- Font Style—Lets you customize the font style.
  - Do not specify—Specifies the default.
  - Use selection—Enables a list of font styles that you can select from.
- Font Size—Lets you customize the font size.
  - Do not specify—Specifies the default.
  - Use selection—Enables a list of font sizes that you can select from.
- Font Weight—Lets you customize the font weight.
  - Do not specify—Specifies the default.
  - Use selection—Enables a list of font weights that you can select from.
- Preview—Displays a preview of your selection.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Select Foreground Color

You can get to this dialog box from various paths.

The Select Foreground Color dialog box lets you conveniently select custom colors that comprise a style.

Fields

- Do not specify—Specifies to use the default.
- Use selection—Enables the Swatches, HSB, and RGB tabs where you can select colors.
  - Swatches tab—Lets you conveniently select custom colors. Click on a color block to select a color.
  - HSB tab—Lets you select custom colors defined by hue, saturation, and brightness (HSB).

Adjust the Spectrum Bar to the basic color of the light spectrum you desire. Then click and drag the mouse over the color plane until you reach the desired shade. Do this for the hue, saturation, and brightness settings. Alternatively, you can adjust the H, S, and B settings manually by clicking the H, S, and B radio buttons and selecting the up or down arrow.

The R, G, and B fields display a translation to RGB values as you drag the mouse.

– RGB tab—Lets you select custom colors defined by red, green, and blue (RGB).

Adjust the red, green, and blue slide bars to the shade that you desire. Alternatively, you can click the up and down arrows on the RGB values to make the values increase or decrease.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Select Background Color

You can get to this dialog box from various paths.

The Select Background Color dialog box lets you conveniently select custom colors that comprise a style.

Fields

- Do not specify—Specifies to use the default.
- Use selection—Enables the Swatches, HSB, and RGB tabs where you can select colors.
  – Swatches tab—Lets you conveniently select custom colors. Click on a color block to select a color.
  – HSB tab—Lets you select custom colors defined by hue, saturation, and brightness (HSB).

  Adjust the Spectrum Bar to the basic color of the light spectrum you desire. Then click and drag the mouse over the color plane until you reach the desired shade. Do this for the hue, saturation, and brightness settings. Alternatively, you can adjust the H, S, and B settings manually by clicking the H, S, and B radio buttons and selecting the up or down arrow.

  The R, G, and B fields display a translation to RGB values as you drag the mouse.
  – RGB tab—Lets you select custom colors defined by red, green, and blue (RGB).

  Adjust the red, green, and blue slide bars to the shade that you desire. Alternatively, you can click the up and down arrows on the RGB values to make the values increase or decrease.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Page Title Tab

Configuration > VPN > WebVPN > Webpage Customization > Add/Edit Customization Objects > Title Page Tab

The Page Title Tab lets you customize the WebVPN page that appears to WebVPN users when they initially connect to the security appliance, including the page style, the title, and the logo.

Fields

This tab contains several Style fields. Use any Cascading Style Sheet (CSS) parameters to define the style. To easily change the font, background color and foreground color, use the Configure button.

You can use a comma-separated RGB value, an HTML color value, or the name of the color if recognized in HTML.

RGB format is 0,0,0, a range of decimal numbers from 0 to 255 for each color (red, green, blue); the comma separated entry indicates the level of intensity of each color to combine with the others.

HTML format is #000000, six digits in hexadecimal format; the first and second represent red, the third and fourth green, and the fifth and sixth represent blue.

For best results, check published HTML and RGB tables. To find tables online, enter RGB in a search engine. For more information about CSS, visit the World Wide Web Consortium (W3C) website at www.w3.org.

- Page Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and appears HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Title—Lets you configure the title of the WebVPN page, including the text and style of the text.
  - Text—Enter the text that you want to appear in the title.
  - Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and appears HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Titlebar Logo—Lets you specify your own custom logo that appears on the WebVPN page.
  - None—No logo appears on the WebVPN page.
  - Default—The Cisco logo appears on the WebVPN page.
  - Custom—Enter the filename of a custom logo, or click the Browse Flash button to browse for a file.
  - Browse Flash—Browse for a custom file.
  - Upload Logo—Displays the Upload Logo dialog box where you can browse for logo files located on the computer running ASDM.

- Sample Preview—Displays the upload logo using your current title and logo settings.
  - Preview in Browser—Displays your current settings in a browser window.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Page Title Tab > Upload Logo

Configuration > VPN > WebVPN > Webpage Customization > Add/Edit Customization Objects > Title Page Tab > Upload Logo Dialog Box

The Upload Logo panel lets you locate a logo file on the computer you are using to run ASDM, and to upload that logo to the security appliance.

Fields

- Local File Path—Displays the path to the logo that you define using the **Browse Local** button.
- Browse Local—Click to browse the fie structure of the computer you are using to run ASDM and locate the logo.
- Flash File System Path—Displays the path to the logo that you define using the **Browse Flash** button.
- Browse Flash—Click to browse the file structure of Flash memory on the security appliance to decide where you want to locate the logo.
- Upload File—Click to display the Browse Flash dialog box. The name of the logo file you have selected appears in the File Name box. Click OK to set this path to flash memory.

You return to the **Upload Logo** panel. Click the **Upload** button to upload the new logo file to Flash memory. This logo now appears in the **Preview** box.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Login Page Tab > Login Box Tab

Configuration > VPN > WebVPN > Webpage Customization > Add/Edit Customization Objects > Login Page Tab > Login Box Tab

The Login Box tab lets you customize the Login box of the WebVPN page that appears to WebVPN users when they initially connect to the security appliance, including the title and the message.

Fields

This tab contains several Style fields. Use any Cascading Style Sheet (CSS) parameters to define the style. To easily change the font, background color and foreground color, use the Configure button.

You can use a comma-separated RGB value, an HTML color value, or the name of the color if recognized in HTML.

RGB format is 0,0,0, a range of decimal numbers from 0 to 255 for each color (red, green, blue); the comma separated entry indicates the level of intensity of each color to combine with the others.

HTML format is #000000, six digits in hexadecimal format; the first and second represent red, the third and fourth green, and the fifth and sixth represent blue.

For best results, check published HTML and RGB tables. To find tables online, enter RGB in a search engine. For more information about CSS, visit the World Wide Web Consortium (W3C) website at www.w3.org.

- Login Title—Lets you specify the text to appear in the Login box title and the style of the login title.
  - Text—Enter the text that you want to appear in the Title of the Login Box.
  - Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Login Message—Lets you specify the message that appears in the Login box, and the style of that message.
  - Text—Enter the text that you want to appear as the message in the Login box.
  - Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Sample Preview—Displays the Login Title and Login Message using your settings.
  - Preview in Browser—Displays the Login Title and Login Message using your current settings in a browser window.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Login Page Tab > Login Prompts Tab

Configuration > VPN > WebVPN > Webpage Customization > Add/Edit Customization Objects > Login Page Tab > Login Prompts Tab

The Login Prompts tab lets you customize the login prompts of the WebVPN page that appears to WebVPN users when they initially connect to the security appliance, including the username, password, and group prompts.

Fields

This tab contains several Style fields. Use any Cascading Style Sheet (CSS) parameters to define the style. To easily change the font, background color and foreground color, use the Configure button.

You can use a comma-separated RGB value, an HTML color value, or the name of the color if recognized in HTML.

RGB format is 0,0,0, a range of decimal numbers from 0 to 255 for each color (red, green, blue); the comma separated entry indicates the level of intensity of each color to combine with the others.

HTML format is #000000, six digits in hexadecimal format; the first and second represent red, the third and fourth green, and the fifth and sixth represent blue.

For best results, check published HTML and RGB tables. To find tables online, enter RGB in a search engine. For more information about CSS, visit the World Wide Web Consortium (W3C) website at www.w3.org.

- Username Prompt—Lets you customize the username prompt, including the text that appears and the style of that text.
  - Text—Enter the text to display for the username prompt.
  - Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Password Prompt—Lets you customize the password prompt, including the text that appears and the style of that text.
  - Text—Enter the text to display for the password prompt.
  - Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Group Prompt—Lets you customize the group prompt, including the text that appears and the style of that text.
  - Text—Enter the text to display for the group prompt.
  - Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure button—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Sample Preview—Displays the Login box using your current username, password, and group prompt settings.

- Preview in Browser—Displays the Login box using your current username, password, and group prompt settings in a browser window.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Login Page Tab > Login Buttons Tab

Configuration > VPN > WebVPN > Webpage Customization > Add/Edit Customization Objects > Login Page Tab > Login Buttons Tab

The Login Buttons tab lets you customize the Login and Clear buttons of the WebVPN page that appears to WebVPN users when they initially connect to the security appliance.

Fields

This tab contains several Style fields. Use any Cascading Style Sheet (CSS) parameters to define the style. To easily change the font, background color and foreground color, use the Configure button.

You can use a comma-separated RGB value, an HTML color value, or the name of the color if recognized in HTML.

RGB format is 0,0,0, a range of decimal numbers from 0 to 255 for each color (red, green, blue); the comma separated entry indicates the level of intensity of each color to combine with the others.

HTML format is #000000, six digits in hexadecimal format; the first and second represent red, the third and fourth green, and the fifth and sixth represent blue.

For best results, check published HTML and RGB tables. To find tables online, enter RGB in a search engine. For more information about CSS, visit the World Wide Web Consortium (W3C) website at www.w3.org.

- Login Button—Lets you customize the Login button, including the text that appears on the button and the style of the button.
  - Text—Enter the text to display on the Login button.
  - Style—Define the style of the Login button with CSS parameters (maximum 256 characters).
  - Configure button—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Clear—Lets you customize the Clear button, including the text that appears on the button and the style of the button.
  - Text—Enter the text to display on the Clear button.
  - Style—Define the style of the Login button with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.

- Sample Preview—Displays the Login and clear buttons using your current settings.
  - Preview in Browser—Displays the Login and Clear buttons using your current settings in a browser window.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Logout Page Tab

Configuration > VPN > WebVPN > Webpage Customization > Add/Edit Customization Objects > Logout Page Tab

The Logout Page tab lets you customize the Logout page that appears to WebVPN users when they Log out of WebVPN service.

Fields

This tab contains several Style fields. Use any Cascading Style Sheet (CSS) parameters to define the style. To easily change the font, background color and foreground color, use the Configure button.

You can use a comma-separated RGB value, an HTML color value, or the name of the color if recognized in HTML.

RGB format is 0,0,0, a range of decimal numbers from 0 to 255 for each color (red, green, blue); the comma separated entry indicates the level of intensity of each color to combine with the others.

HTML format is #000000, six digits in hexadecimal format; the first and second represent red, the third and fourth green, and the fifth and sixth represent blue.

For best results, check published HTML and RGB tables. To find tables online, enter RGB in a search engine. For more information about CSS, visit the World Wide Web Consortium (W3C) website at www.w3.org.

- Logout Title—Lets you specify the text to appear in the Logout box title and the style of the logout title.
  - Text—Enter the text that you want to appear as the message in the Logout page.
  - Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Logout Message—Lets you specify the message to appear on the Logout page.
  - Text—Enter the text that you want to appear as the message in the Logout page.
  - Style—Define the style with CSS parameters (maximum 256 characters).

  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Sample Preview—Displays the Logout page using your settings.
  - Preview in Browser—Displays the page using your current settings in a browser window.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Home Page Tab > Border Color Tab

Configuration > VPN > WebVPN > Webpage Customization > Add/Edit Customization Objects > Home Page Tab > Border Color Tab

The Border Color tab lets you customize the border of the WebVPN Home page that appears to WebVPN users after they are authenticated by the security appliance.

Fields

- Border Style—Use any Cascading Style Sheet (CSS) parameters to define the style, including font styles, and HTML and RGB colors. To easily change the style, use the **Configure** button.

  You can use a comma-separated RGB value, an HTML color value, or the name of the color if recognized in HTML.

  RGB format is 0,0,0, a range of decimal numbers from 0 to 255 for each color (red, green, blue); the comma separated entry indicates the level of intensity of each color to combine with the others.

  HTML format is #000000, six digits in hexadecimal format; the first and second represent red, the third and fourth green, and the fifth and sixth represent blue.

  For best results, check published HTML and RGB tables. To find tables online, enter RGB in a search engine. For more information about CSS, visit the World Wide Web Consortium (W3C) website at www.w3.org.
- Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Sample Preview—Displays the WebVPN Home page using your border settings.
  - Preview in Browser—Displays the WebVPN Home page using your border settings in a browser window.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Home Page Tab > Web Applications Tab

Configuration > VPN > WebVPN > Webpage Customization > Add/Edit Customization Objects > Home Page Tab > Web Applications Tab

The Web Applications tab lets you customize the Web Applications box of the WebVPN Home page that appears to authenticated WebVPN users.

Fields

This tab contains several Style fields. Use any Cascading Style Sheet (CSS) parameters to define the style. To easily change the font, background color and foreground color, use the Configure button.

You can use a comma-separated RGB value, an HTML color value, or the name of the color if recognized in HTML.

RGB format is 0,0,0, a range of decimal numbers from 0 to 255 for each color (red, green, blue); the comma separated entry indicates the level of intensity of each color to combine with the others.

HTML format is #000000, six digits in hexadecimal format; the first and second represent red, the third and fourth green, and the fifth and sixth represent blue.

For best results, check published HTML and RGB tables. To find tables online, enter RGB in a search engine. For more information about CSS, visit the World Wide Web Consortium (W3C) website at www.w3.org.

- Title—Lets you customize the title of the Web Applications box.
  - Text—Enter the text that you want to appear as the title.
  - Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Message—Lets you customize the message (under the title) of the Web Applications box,
  - Text—Enter the text that you want to appear as the message.
  - Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Dropdown—Lets you customize the drop-down list of the Web Applications box.
  - Text—Enter the text that you want to appear in the drop-down list.
  - Style—Define the style with CSS parameters (maximum 256 characters).

  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Sample Preview—Displays the WebVPN Home page using your Web Application settings.
  - Preview in Browser—Displays the WebVPN Home page using your Web Application in a browser window.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Home Page Tab > Application Access Tab

Configuration > VPN > WebVPN > Webpage Customization > Add/Edit Customization Objects > Home Page Tab > Application Access Tab

The Applications Access tab lets you customize the Applications Access box of the WebVPN Home page that appears to authenticated WebVPN users.

Fields

This tab contains several Style fields. Use any Cascading Style Sheet (CSS) parameters to define the style. To easily change the font, background color and foreground color, use the Configure button.

You can use a comma-separated RGB value, an HTML color value, or the name of the color if recognized in HTML.

RGB format is 0,0,0, a range of decimal numbers from 0 to 255 for each color (red, green, blue); the comma separated entry indicates the level of intensity of each color to combine with the others.

HTML format is #000000, six digits in hexadecimal format; the first and second represent red, the third and fourth green, and the fifth and sixth represent blue.

For best results, check published HTML and RGB tables. To find tables online, enter RGB in a search engine. For more information about CSS, visit the World Wide Web Consortium (W3C) website at www.w3.org.

- Title—Lets you customize the title of the Applications Access box.
  - Text—Enter the text that you want to appear as the title.
  - Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Message—Lets you customize the message under the title of the Applications Access box
  - Text—Enter the text that you want to appear as the message.

  - Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Sample Preview—Displays the WebVPN Home page using your Application Access settings.
  - Preview in Browser—Displays the WebVPN Home page using your Application Access settings in a browser window.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Home Page Tab > Browse Network Tab

Configuration > VPN > WebVPN > Webpage Customization > Add/Edit Customization Objects > Home Page Tab > Browse Network Tab

The Browse Networks tab lets you customize the Browse Networks box of the WebVPN Home page that appears to authenticated WebVPN users.

Fields

This tab contains several Style fields. Use any Cascading Style Sheet (CSS) parameters to define the style. To easily change the font, background color and foreground color, use the Configure button.

You can use a comma-separated RGB value, an HTML color value, or the name of the color if recognized in HTML.

RGB format is 0,0,0, a range of decimal numbers from 0 to 255 for each color (red, green, blue); the comma separated entry indicates the level of intensity of each color to combine with the others.

HTML format is #000000, six digits in hexadecimal format; the first and second represent red, the third and fourth green, and the fifth and sixth represent blue.

For best results, check published HTML and RGB tables. To find tables online, enter RGB in a search engine. For more information about CSS, visit the World Wide Web Consortium (W3C) website at www.w3.org.

- Title—Lets you customize the title of the Browse Networks box.
  - Text—Enter the text that you want to appear as the title.
  - Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Message—Lets you customize the message under the title of the Browse Networks box

  - Text—Enter the text that you want to appear as the message.
  - Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Dropdown—Lets you customize the drop-down list of the Browse Networks box.
  - Text—Enter the text that you want to appear in the drop-down list.
  - Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Sample Preview—Displays the WebVPN Home page using your Browse Networks settings.
  - Preview in Browser—Displays the WebVPN Home page using your Browse Networks settings in a browser window.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Home Page Tab > Web Bookmarks Tab

Configuration > VPN > WebVPN > Webpage Customization > Add/Edit Customization Objects > Home Page Tab > Web Bookmarks Tab

The Web Bookmarks tab lets you customize the Web Bookmarks title and the appearance of the bookmarks links on the WebVPN Home page that appears to authenticated WebVPN users.

Fields

This tab contains several Style fields. Use any Cascading Style Sheet (CSS) parameters to define the style. To easily change the font, background color and foreground color, use the Configure button.

You can use a comma-separated RGB value, an HTML color value, or the name of the color if recognized in HTML.

RGB format is 0,0,0, a range of decimal numbers from 0 to 255 for each color (red, green, blue); the comma separated entry indicates the level of intensity of each color to combine with the others.

HTML format is #000000, six digits in hexadecimal format; the first and second represent red, the third and fourth green, and the fifth and sixth represent blue.

For best results, check published HTML and RGB tables. To find tables online, enter RGB in a search engine. For more information about CSS, visit the World Wide Web Consortium (W3C) website at www.w3.org.

- Bookmark Title—Lets you customize the title.
  - Text—Enter the text that you want to appear as the title.
  - Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Bookmark Links Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools
- Sample Preview—Displays the WebVPN Home page using your Web Bookmarks settings.
  - Preview in Browser—Displays the WebVPN Home page using your Web Bookmarks settings in a browser window.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Home Page Tab > File Bookmarks Tab

Configuration > VPN > WebVPN > Webpage Customization > Add/Edit Customization Objects > Home Page Tab > File Bookmarks Tab

The File Bookmarks tab lets you customize the File Bookmarks title and the appearance of the bookmarks links on the WebVPN Home page that appears to authenticated WebVPN users.

Fields

This tab contains several Style fields. Use any Cascading Style Sheet (CSS) parameters to define the style. To easily change the font, background color and foreground color, use the Configure button.

You can use a comma-separated RGB value, an HTML color value, or the name of the color if recognized in HTML.

RGB format is 0,0,0, a range of decimal numbers from 0 to 255 for each color (red, green, blue); the comma separated entry indicates the level of intensity of each color to combine with the others.

HTML format is #000000, six digits in hexadecimal format; the first and second represent red, the third and fourth green, and the fifth and sixth represent blue.

For best results, check published HTML and RGB tables. To find tables online, enter RGB in a search engine. For more information about CSS, visit the World Wide Web Consortium (W3C) website at www.w3.org.

- Bookmark Title—Lets you customize the title.

  - Text—Enter the text that you want to appear as the title.
  - Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Bookmark Links Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools
- Sample Preview—Displays the WebVPN Home page using your File Bookmarks settings.
  - Preview in Browser—Displays the WebVPN Home page using your File Bookmarks settings in a browser window.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Application Access Window Tab

Configuration > VPN > WebVPN > Webpage Customization > Add/Edit Customization Objects > Application Access Window Tab

The Application Access Window tab lets you customize the Application Access window that appears to authenticated WebVPN users that select Application Access on the WebVPN Home page.

Fields

This tab contains several Style fields. Use any Cascading Style Sheet (CSS) parameters to define the style. To easily change the font, background color and foreground color, use the Configure button.

You can use a comma-separated RGB value, an HTML color value, or the name of the color if recognized in HTML.

RGB format is 0,0,0, a range of decimal numbers from 0 to 255 for each color (red, green, blue); the comma separated entry indicates the level of intensity of each color to combine with the others.

HTML format is #000000, six digits in hexadecimal format; the first and second represent red, the third and fourth green, and the fifth and sixth represent blue.

For best results, check published HTML and RGB tables. To find tables online, enter RGB in a search engine. For more information about CSS, visit the World Wide Web Consortium (W3C) website at www.w3.org.

- Window Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.

- Warning Message—Enter the text that you want to appear as the warning message.
- Show application details in the application access window—Lets you disable the display of application details that appear on the Application Access Window.
- Sample Preview—Displays the Application Access Window using your settings.
  - Preview in Browser—Displays the Application Access Window using your settings in a browser window.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Prompt Dialog Tab

Configuration > VPN > WebVPN > Webpage Customization > Add/Edit Customization Objects > Prompt Dialog Tab

The Prompt Dialog tab lets you customize the appearance of dialog messages that appear to authenticated WebVPN users.

Fields

This tab contains several Style fields. Use any Cascading Style Sheet (CSS) parameters to define the style. To easily change the font, background color and foreground color, use the Configure button.

You can use a comma-separated RGB value, an HTML color value, or the name of the color if recognized in HTML.

RGB format is 0,0,0, a range of decimal numbers from 0 to 255 for each color (red, green, blue); the comma separated entry indicates the level of intensity of each color to combine with the others.

HTML format is #000000, six digits in hexadecimal format; the first and second represent red, the third and fourth green, and the fifth and sixth represent blue.

For best results, check published HTML and RGB tables. To find tables online, enter RGB in a search engine. For more information about CSS, visit the World Wide Web Consortium (W3C) website at www.w3.org.

- Dialog Title Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Dialog Message Style—Define the style with CSS parameters (maximum 256 characters).
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Dialog Border Style—Define the style with CSS parameters (maximum 256 characters).

  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
- Sample Preview—Displays a sample of a dialog message using your settings.
  - Preview in Browser—Displays a sample of a dialog message using your settings in a browser window.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Add/Edit Webpage Customization Object > Quick Style Configuration

Configuration > VPN > WebVPN > Webpage Customization > Add/Edit Customization Objects > Quick Style Configuration >

The Quick Style Configuration dialog box lets you apply a single style to multiple WebVPN window customization settings.

Fields

- Select Fields—click the fields that you want to share a single style.
- Specify Style—Lets you specify a custom style to use for the fields you have selected.
  - Use custom style—click disable the default style, and supply the custom style.
  - Style—Define the style of the WebVPN page with Cascading Style Sheet (CSS) parameters (maximum 256 characters), including font styles, and HTML and RGB colors. To easily change the style, use the **Configure** button.

    You can use a comma-separated RGB value, an HTML color value, or the name of the color if recognized in HTML.

    RGB format is 0,0,0, a range of decimal numbers from 0 to 255 for each color (red, green, blue); the comma separated entry indicates the level of intensity of each color to combine with the others.

    HTML format is #000000, six digits in hexadecimal format; the first and second represent red, the third and fourth green, and the fifth and sixth represent blue.

    For best results, check published HTML and RGB tables. To find tables online, enter RGB in a search engine. For more information about CSS, visit the World Wide Web Consortium (W3C) website at www.w3.org.
  - Configure—Lets you configure the font, foreground color, and background color, and displays HTML color swatches, and HSB (Hue, Saturation, Brightness) and RGB (Red, Green, Blue) selection tools.
  - Preview—Displays a sample of the style you selected.
- Use default styles—Enables default styles.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

CHAPTER 30

WebVPN End User Set-up

This ections is for the system administrator who sets up WebVPN for end users. It summarizes configuration requirements and tasks for the user remote system. It also specifies information to communicate to users to get them started using WebVPN. This section includes the following topics:

- Requiring Usernames and Passwords
- Communicating Security Tips
- Configuring Remote Systems to Use WebVPN Features
- Capturing WebVPN Data

Note We assume you have already configured the security appliance for WebVPN.

Requiring Usernames and Passwords

Depending on your network, during a remote session users might have to log in to any or all of the following: the computer itself, an Internet service provider, WebVPN, mail or file servers, or corporate applications. Users might have to authenticate in many different contexts, requiring different information, such as a unique username, password, or PIN.

Table 30-1 lists the type of usernames and passwords that WebVPN users might need to know.

Table 30-1 Usernames and Passwords to Give to WebVPN Users

Login Username/ Password Type	Purpose	Entered When
Computer	Access the computer	Starting the computer
Internet Service Provider	Access the Internet	Connecting to an Internet service provider
WebVPN	Access remote network	Starting WebVPN
File Server	Access remote file server	Using the WebVPN file browsing feature to access a remote file server
Corporate Application Login	Access firewall-protected internal server	Using the WebVPN web browsing feature to access an internal protected website
Mail Server	Access remote mail server via WebVPN	Sending or receiving e-mail messages

Communicating Security Tips

Advise users always to log out from the WebVPN session. (To log out of WebVPN, click the logout icon on the WebVPN toolbar or close the browser.)

Advise users that using WebVPN does not ensure that communication with every site is secure. WebVPN ensures the security of data transmission between the remote PC or workstation and the security appliance on the corporate network. If a user then accesses a non-HTTPS web resource (located on the Internet or on the internal network), the communication from the corporate security appliance to the destination web server is not secure.

Configuring Remote Systems to Use WebVPN Features

Table 30-2 includes the following information about setting up remote systems to use WebVPN:

- Starting WebVPN
- Using the WebVPN Floating Toolbar
- Web Browsing
- Network Browsing and File Management
- Using Applications (Port Forwarding)
- Using E-mail via Port Forwarding
- Using E-mail via Web Access
- Using E-mail via e-mail proxy

Table 30-2 also provides information about the following:

- WebVPN requirements, by feature
- WebVPN supported applications
- Client application installation and configuration requirements
- Information you might need to provide end users
- Tips and use suggestions for end users

It is possible you have configured user accounts differently and that different WebVPN features are available to each user. Table 30-2 organizes information by feature, so you can skip over the information for unavailable features.

Table 30-2 WebVPN Remote System Configuration and End User Requirements

Task	Remote System or End User Requirements	Specifications or Use Suggestions
Starting WebVPN	Connection to the Internet	Any Internet connection is supported, including: • Home DSL, cable, or dial-ups • Public kiosks • Hotel hook-ups • Airport wireless nodes • Internet cafes
	WebVPN-supported browser	We recommend the following browsers for WebVPN. Other browsers might not fully support WebVPN features. On Microsoft Windows: • Internet Explorer version 6.0 • Netscape version 7.2 • Mozilla version 1.7 and later • Firefox 1.x On Linux: • Mozilla version 1.7 • Netscape version 7.2 • Firefox 1.x On Solaris: • Netscape version 7.2 On Macintosh OS X: • Safari version 1.0 • Firefox 1.x
	Cookies enabled on browser	Cookies must be enabled on the browser in order to access applications via port forwarding.
	URL for WebVPN	An https address in the following form: https://*address* where *address* is the IP address or DNS hostname of an interface of the security appliance (or load balancing cluster) on which WebVPN is enabled. For example: https://10.89.192.163 or https://cisco.example.com.
	WebVPN username and password	
	[Optional] Local printer	WebVPN does not support printing from a web browser to a network printer. Printing to a local printer is supported.

Table 30-2 WebVPN Remote System Configuration and End User Requirements (continued)

Task	Remote System or End User Requirements	Specifications or Use Suggestions
Using the WebVPN Floating Toolbar		A floating toolbar is available to simplify the use of WebVPN. The toolbar lets you enter URLs, browse file locations, and choose preconfigured web connections without interfering with the main browser window. If you configure your browser to block popups, the floating toolbar cannot display. The floating toolbar represents the current WebVPN session. If you click the **Close** button, the security appliance prompts you to confirm that you want to close the WebVPN session. **Tip** TIP: To paste text into a text field, use Ctrl-V. (Right-clicking is disabled on the WebVPN toolbar.)
Web Browsing	Usernames and passwords for protected websites	Using WebVPN does not ensure that communication with every site is secure. See "Communicating Security Tips." The look and feel of web browsing with WebVPN might be different from what users are accustomed to. For example: • The WebVPN title bar appears above each web page • You access websites by: – Entering the URL in the Enter Web Address field on the WebVPN Home page – Clicking on a preconfigured website link on the WebVPN Home page – Clicking a link on a webpage accessed via one of the previous two methods Also, depending on how you configured a particular account, it might be that: • Some websites are blocked • Only the websites that appear as links on the WebVPN Home page are available

Table 30-2 WebVPN Remote System Configuration and End User Requirements (continued)

Task	Remote System or End User Requirements	Specifications or Use Suggestions
Network Browsing and File Management	File permissions configured for shared remote access	Only shared folders and files are accessible via WebVPN.
	Server name and passwords for protected file servers	—
	Domain, workgroup, and server names where folders and files reside	Users might not be familiar with how to locate their files through your organization network.
	—	Do not interrupt the **Copy File to Server** command or navigate to a different screen while the copying is in progress. Interrupting the operation can cause an incomplete file to be saved on the server.

Table 30-2 WebVPN Remote System Configuration and End User Requirements (continued)

Task	Remote System or End User Requirements	Specifications or Use Suggestions
Using Applications	**Note** On Macintosh OS X, only the Safari browser supports this feature.	
(called Port Forwarding or Application Access)	**Note** Because this feature requires installing Sun Microsystems Java™ Runtime Environment and configuring the local clients, and because doing so requires administrator permissions on the local system, it is unlikely that users will be able to use applications when they connect from public remote systems.	
	⚠ **Caution** Users should always close the Application Access window when they finish using applications by clicking the **Close** icon. Failure to quit the window properly can cause Application Access or the applications themselves to be disabled. See	
	Client applications installed	—
	Cookies enabled on browser	—
	Administrator privileges	User must have administrator access on the PC if you use DNS names to specify servers because modifying the hosts file requires it.
	Sun Microsystems Java Runtime Environment (JRE) version 1.4.x and 1.5.x installed. Javascript must be enabled on the browser. By default, it is enabled.	If JRE is not installed, a pop-up window displays, directing users to a site where it is available. On rare occasions, the WebVPN port forwarding applet fails with JAVA exception errors. If this happens, do the following: 1. Clear the browser cache and close the browser. 2. Verify that no JAVA icons are in the computer task bar. Close all instances of JAVA. 3. Establish a WebVPN session and launch the port forwarding JAVA applet.
	Client applications configured, if necessary. **Note** The Microsoft Outlook client does not require this configuration step. All non-Windows client applications require configuration. To see if configuration is necessary for a Windows application, check the value of the Remote Server. • If the Remote Server contains the server hostname, you do not need to configure the client application. • If the Remote Server field contains an IP address, you must configure the client application.	To configure the client application, use the server's locally mapped IP address and port number. To find this information: 1. Start WebVPN on the remote system and click the Application Access link on the WebVPN Home page. The Application Access window appears. 2. In the Name column, find the name of the server you want to use, then identify its corresponding client IP address and port number (in the Local column). 3. Use this IP address and port number to configure the client application. Configuration steps vary for each client application.
	Note Clicking a URL (such as one in an -e-mail message) in an application running over WebVPN does not open the site over WebVPN. To open a site over WebVPN, cut and paste the URL into the Enter WebVPN (URL) Address field.	

Table 30-2 WebVPN Remote System Configuration and End User Requirements (continued)

Task	Remote System or End User Requirements	Specifications or Use Suggestions
Using E-mail via Application Access	Fulfill requirements for Application Access (See Using Applications)	To use mail, start Application Access from the WebVPN Home page. The mail client is then available for use.
	Note If you are using an IMAP client and you lose your mail server connection or are unable to make a new connection, close the IMAP application and restart WebVPN.	
	Other mail clients	We have tested Microsoft Outlook Express versions 5.5 and 6.0. WebVPN should support other SMTPS, POP3S, or IMAP4S e-mail programs via port forwarding, such as Netscape Mail, Lotus Notes, and Eudora, but we have not verified them.
Using E-mail via Web Access	Web-based e-mail product installed	Supported products include: • Outlook Web Access For best results, use OWA on Internet Explorer 6.x or higher, Mozilla 1.7, or Firefox 1.x. • Lotus iNotes Other web-based e-mail products should also work, but we have not verified them.
Using E-mail via E-mail Proxy	SSL-enabled mail application installed Do not set the security appliance SSL version to TLSv1 Only. Outlook and Outlook Express do not support TLS.	Supported mail applications: • Microsoft Outlook • Microsoft Outlook Express versions 5.5 and 6.0 • Netscape Mail version 7 • Eudora 4.2 for Windows 2000 Other SSL-enabled mail clients should also work, but we have not verified them.
	Mail application configured	

Capturing WebVPN Data

The CLI capture command lets you log information about websites that do not display properly over a WebVPN connection. This data can help your Cisco customer support engineer troubleshoot problems. The following sections describe how to use the capture command:

- Creating a Capture File
- Using a Browser to Display Capture Data

Note Enabling WebVPN capture affects the performance of the security appliance. Be sure to disable the capture after you generate the capture files needed for troubleshooting.

Creating a Capture File

Perform the following steps to capture data about a WebVPN session to a file.

Step 1 To start the WebVPN capture utility, use the **capture** command from privileged EXEC mode.

capture *capture_name* **type webvpn user** *webvpn_username*

where:

- *capture_name* is a name you assign to the capture, which is also prepended to the name of the capture files.
- *webvpn_user* is the username to match for capture.

The capture utility starts.

Step 2 A WebVPN user logs in to begin a WebVPN session. The capture utility is capturing packets.

Stop the capture by using the **no** version of the command.

no capture *capture_name*

The capture utility creates a *capture_name*.zip file, which is encrypted with the password **koleso**.

Step 3 Send the .zip file to Cisco Systems, or attach it to a Cisco TAC service request.

Step 4 To look at the contents of the .zip file, unzip it using the password **koleso**.

The following example creates a capture named *hr*, which captures WebVPN traffic for user2 to a file:

```
hostname# capture hr type webvpn user user2
WebVPN capture started.
   capture name   hr
   user name      user2
hostname# no capture hr
```

Using a Browser to Display Capture Data

Perform the following steps to capture data about a WebVPN session and view it in a browser.

Step 1 To start the WebVPN capture utility, use the **capture** command from privileged EXEC mode.

capture *capture_name* **type webvpn user** *webvpn_username*

where:

- *capture_name* is a name you assign to the capture, which is also prepended to the name of the capture files.
- *webvpn_username* is the username to match for capture.

The capture utility starts.

Step 2 A WebVPN user logs in to begin a WebVPN session. The capture utility is capturing packets.

Stop the capture by using the **no** version of the command.

Step 3 Open a browser and in the address box enter

https://*IP_address or hostname of the security appliance***/webvpn_capture.html**

The captured content displays in a sniffer format.

Step 4 When you finish examining the capture content, stop the capture by using the **no** version of the command.

CHAPTER 31

E-Mail Proxy

E-mail proxies extend remote e-mail capability to WebVPN users. When users attempt an e-mail session via e-mail proxy, the e-mail client establishes a tunnel using the SSL protocol.

The e-mail proxy protocols are as follows:

POP3S

POP3S is one of the e-mail proxies WebVPN supports. By default the Security Appliance listens to port 995, and connections are automatically allowed to port 995 or to the configured port. The POP3 proxy allows only SSL connections on that port. After the SSL tunnel establishes, the POP3 protocol starts, and then authentication occurs. POP3S is for receiving e-mail.

IMAP4S

IMAP4S is one of the e-mail proxies WebVPN supports. By default the Security Appliance listens to port 993, and connections are automatically allowed to port 993 or to the configured port. The IMAP4 proxy allows only SSL connections on that port. After the SSL tunnel establishes, the IMAP4 protocol starts, and then authentication occurs. IMAP4S is for receiving e-mail.

SMTPS

SMTPS is one of the e-mail proxies WebVPN supports. By default the Security Appliance listens to port 988, and connections are automatically allowed to port 988 or to the configured port. The SMTPS proxy allows only SSL connections on that port. After the SSL tunnel establishes, the SMTPS protocol starts, and then authentication occurs. SMTPS is for sending e-mail.

Configuring E-Mail Proxy

Configuring e-mail proxy on the consists of the following tasks:

- Enabling e-Mail proxy on interfaces.
- Configuring e-mail proxy default servers.
- Setting AAA server groups and a default group policy.
- Configuring delimiters.

Configuring E-mail proxy also has these requirements:

- Users who access e-mail from both local and remote locations via e-mail proxy require separate e-mail accounts on their e-mail program for local and remote access.
- E-mail proxy sessions require that the user authenticate.

AAA

Configuration > Features > VPN > E-mail Proxy > AAA

This panel has three tabs:

- POP3S Tab
- IMAP4S Tab
- SMTPS Tab

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

POP3S Tab

Configuration > Features > VPN > E-mail Proxy > AAA > POP3S Tab

The POP3S AAA panel associates AAA server groups and configures the default group policy for POP3S sessions.

Fields

- AAA server groups—Click to go to the AAA Server Groups panel (Configuration > Features > Properties > AAA Setup > AAA Server Groups), where you can add or edit AAA server groups.
- group policies—Click to go to the Group Policy panel (Configuration > Features > VPN > General > Group Policy), where you can add or edit group policies.
- Authentication Server Group—Select the authentication server group for POP3S user authentication. The default is to have no authentication servers configured. If you have set AAA as the authentication method for POP3S (Configuration > Features AAA > VPN > E-Mail Proxy > Authentication panel), you must configure an AAA server and select it here, or authentication always fails.
- Authorization Server Group—Select the authorization server group for POP3S user authorization. The default is to have no authorization servers configured.
- Accounting Server Group—Select the accounting server group for POP3S user accounting. The default is to have no accounting servers configured.
- Default Group Policy—Select the group policy to apply to POP3S users when AAA does not return a CLASSID attribute. The length must be between 4 and 15 alphanumeric characters. If you do not specify a default group policy, and there is no CLASSID, the security appliance can not establish the session.
- Authorization Settings—Lets you set values for usernames that the security appliance recognizes for POP3S authorization. This applies to POP3S users that authenticate with digital certificates and require LDAP or RADIUS authorization.

- User the entire DN as the username—Select to use the Distinguished Name for POP3S authorization.
- Specify individual DN fields as the username—Select to specify specific DN fields for user authorization.

  You can choose two DN fields, primary and secondary. For example, if you choose EA, users authenticate according to their e-mail address. Then a user with the Common Name (CN) John Doe and an e-mail address of johndoe@cisco.com cannot authenticate as John Doe or as johndoe. He must authenticate as johndoe@cisco.com. If you choose EA and O, John Does must authenticate as johndoe@cisco.com and Cisco Systems, Inc.
- Primary DN Field—Select the primary DN field you want to configure for POP3S authorization. The default is CN. Options include the following:

DN Field	Definition
Country (C)	The two-letter country abbreviation. These codes conform to ISO 3166 country abbreviations.
Common Name (CN)	The name of a person, system, or other entity. This is the lowest (most specific) level in the identification hierarchy.
DN Qualifier (DNQ)	A specific DN attribute.
E-mail Address (EA)	The e-mail address of the person, system or entity that owns the certificate.
Generational Qualifier (GENQ)	A generational qualifier such as Jr., Sr., or III.
Given Name (GN)	The first name of the certificate owner.
Initials (I)	The first letters of each part of the certificate owner's name.
Locality (L)	The city or town where the organization is located.
Name (N)	The name of the certificate owner.
Organization (O)	The name of the company, institution, agency, association, or other entity.
Organizational Unit (OU)	The subgroup within the organization.
Serial Number (SER)	The serial number of the certificate.
Surname (SN)	The family name or last name of the certificate owner.
State/Province (S/P)	The state or province where the organization is located.
Title (T)	The title of the certificate owner, such as Dr.
User ID (UID)	The identification number of the certificate owner.

- Secondary DN Field—(Optional) Select the secondary DN field you want to configure for POP3S authorization. The default is OU. Options include all of those in the preceding table, with the addition of **None**, which you select if you do not want to include a secondary field.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

IMAP4S Tab

Configuration > Features > VPN > E-mail Proxy > AAA > IMAP4S Tab

The IMAP4S AAA panel associates AAA server groups and configures the default group policy for IMAP4S sessions.

Fields

- AAA server groups—Click to go to the AAA Server Groups panel (Configuration > Features > Properties > AAA Setup > AAA Server Groups), where you can add or edit AAA server groups.
- group policy—Click to go to the Group Policy panel (Configuration > Features > VPN > General > Group Policy), where you can add or edit group policies.
- Authentication Server Group—Select the authentication server group for IMAP4S user authentication. The default is to have no authentication servers configured. If you have set AAA as the authentication method for IMAP4S (Configuration > Features AAA > VPN > E-Mail Proxy > Authentication panel), you must configure an AAA server and select it here, or authentication always fails.
- Authorization Server Group—Select the authorization server group for IMAP4S user authorization. The default is to have no authorization servers configured.
- Accounting Server Group—Select the accounting server group for IMAP4S user accounting. The default is to have no accounting servers configured.
- Default Group Policy—Select the group policy to apply to IMAP4S users when AAA does not return a CLASSID attribute. If you do not specify a default group policy, and there is no CLASSID, the security appliance can not establish the session.
- Authorization Settings—Lets you set values for usernames that the security appliance recognizes for IMAP4S authorization. This applies to IMAP4S users that authenticate with digital certificates and require LDAP or RADIUS authorization.
  - User the entire DN as the username—Select to use the fully qualified domain name for IMAP4S authorization.
  - Specify individual DN fields as the username—Select to specify specific DN fields for user authorization.

    You can choose two DN fields, primary and secondary. For example, if you choose EA, users authenticate according to their e-mail address. Then a user with the Common Name (CN) John Doe and an e-mail address of johndoe@cisco.com cannot authenticate as John Doe or as johndoe. He must authenticate as johndoe@cisco.com. If you choose EA and O, John Does must authenticate as johndoe@cisco.com *and* Cisco. Systems, Inc.

– **Primary DN Field**—Select the primary DN field you want to configure for IMAP4S authorization. The default is CN. Options include the following:

DN Field	Definition
Country (C)	The two-letter country abbreviation. These codes conform to ISO 3166 country abbreviations.
Common Name (CN)	The name of a person, system, or other entity. This is the lowest (most specific) level in the identification hierarchy.
DN Qualifier (DNQ)	A specific DN attribute.
E-mail Address (EA)	The e-mail address of the person, system or entity that owns the certificate.
Generational Qualifier (GENQ)	A generational qualifier such as Jr., Sr., or III.
Given Name (GN)	The first name of the certificate owner.
Initials (I)	The first letters of each part of the certificate owner's name.
Locality (L)	The city or town where the organization is located.
Name (N)	The name of the certificate owner.
Organization (O)	The name of the company, institution, agency, association, or other entity.
Organizational Unit (OU)	The subgroup within the organization.
Serial Number (SER)	The serial number of the certificate.
Surname (SN)	The family name or last name of the certificate owner.
State/Province (S/P)	The state or province where the organization is located.
Title (T)	The title of the certificate owner, such as Dr.
User ID (UID)	The identification number of the certificate owner.

– Secondary DN Field—(Optional) Select the secondary DN field you want to configure for IMAP4S authorization. The default is OU. Options include all of those in the preceding table, with the addition of None, which you select if you do not want to include a secondary field.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

SMTPS Tab

Configuration > Features > VPN > E-mail Proxy > AAA > SMTPS Tab

The SMTPS AAA panel associates AAA server groups and configures the default group policy for SMTPS sessions.

Fields

- AAA server groups—Click to go to the AAA Server Groups panel (Configuration > Features > Properties > AAA Setup > AAA Server Groups), where you can add or edit AAA server groups.
- group policy—Click to go to the Group Policy panel (Configuration > Features > VPN > General > Group Policy), where you can add or edit group policies.
- Authentication Server Group—Select the authentication server group for SMTPS user authentication. The default is to have no authentication servers configured. If you have set AAA as the authentication method for SMTPS (Configuration > Features AAA > VPN > E-Mail Proxy > Authentication panel), you must configure an AAA server and select it here, or authentication always fails.
- Authorization Server Group—Select the authorization server group for SMTPS user authorization. The default is to have no authorization servers configured.
- Accounting Server Group—Select the accounting server group for SMTPS user accounting. The default is to have no accounting servers configured.
- Default Group Policy—Select the group policy to apply to SMTPS users when AAA does not return a CLASSID attribute. If you do not specify a default group policy, and there is no CLASSID, the security appliance can not establish the session.
- Authorization Settings—Lets you set values for usernames that the security appliance recognizes for SMTPS authorization. This applies to SMTPS users that authenticate with digital certificates and require LDAP or RADIUS authorization.
  - User the entire DN as the username—Select to use the fully qualified domain name for SMTPS authorization.
  - Specify individual DN fields as the username—Select to specify specific DN fields for user authorization.

    You can choose two DN fields, primary and secondary. For example, if you choose EA, users authenticate according to their e-mail address. Then a user with the Common Name (CN) John Doe and an e-mail address of johndoe@cisco.com cannot authenticate as John Doe or as johndoe. He must authenticate as johndoe@cisco.com. If you choose EA and O, John Does must authenticate as johndoe@cisco.com *and* Cisco. Systems, Inc.
  - Primary DN Field—Select the primary DN field you want to configure for SMTPS authorization. The default is CN. Options include the following:

DN Field	Definition
Country (C)	The two-letter country abbreviation. These codes conform to ISO 3166 country abbreviations.
Common Name (CN)	The name of a person, system, or other entity. This is the lowest (most specific) level in the identification hierarchy.
DN Qualifier (DNQ)	A specific DN attribute.
E-mail Address (EA)	The e-mail address of the person, system or entity that owns the certificate.
Generational Qualifier (GENQ)	A generational qualifier such as Jr., Sr., or III.
Given Name (GN)	The first name of the certificate owner.
Initials (I)	The first letters of each part of the certificate owner's name.
Locality (L)	The city or town where the organization is located.
Name (N)	The name of the certificate owner.

DN Field	Definition
Organization (O)	The name of the company, institution, agency, association, or other entity.
Organizational Unit (OU)	The subgroup within the organization.
Serial Number (SER)	The serial number of the certificate.
Surname (SN)	The family name or last name of the certificate owner.
State/Province (S/P)	The state or province where the organization is located.
Title (T)	The title of the certificate owner, such as Dr.
User ID (UID)	The identification number of the certificate owner.

– **Secondary DN Field**—(Optional) Select the secondary DN field you want to configure for SMTPS authorization. The default is OU. Options include all of those in the preceding table, with the addition of None, which you select if you do not want to include a secondary field.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Access

Configuration > VPN > E-Mail Proxy > Access

The E-mail Proxy Access screen lets you identify interfaces on which to configure e-mail proxy. You can configure e-mail proxies on individual interfaces, and you can configure e-mail proxies for one interface and then apply your settings to all interfaces. You cannot configure e-mail proxies for management-only interfaces, or for subinterfaces.

Fields

- Interface—Displays the names of all configured interfaces.
- POP3S Enabled—Shows whether POP3S is enabled for the interface.
- IMAP4s Enabled—Shows whether IMAP4S is enabled for the interface.
- SMTPS Enabled—Shows whether SMTPS is enabled for the interface.
- Edit—Click to edit the e-mail proxy settings for the highlighted interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Edit E-Mail Proxy Access

Configuration > VPN > E-Mail Proxy > Access > Edit E-Mail Proxy Access

The E-mail Proxy Access screen lets you identify interfaces on which to configure e-mail proxy. You can configure e-mail proxies on individual interfaces, and you can configure e-mail proxies for one interface and then apply your settings to all interfaces.

Fields

- Interface—Displays the name of the selected interface.
- POP3S Enabled—Select to enable POP3S for the interface.
- IMAP4s Enabled—elect to enable IMAP4S for the interface.
- SMTPS Enabled—Select to enable SMTPS for the interface.
- Apply to all interface—Select to apply the settings for the current interface to all configured interfaces.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Authentication

Configuration > Features > VPN >E-mail Proxy > Authentication

This panel lets you configure authentication methods for e-mail proxy sessions.

Fields

POP3S/IMAP4S/SMTPS Authentication—Let you configure authentication methods for each of the e-mail proxy types. You can select multiple methods of authentication.

- AAA—Select to require AAA authentication. This option requires a configured AAA server. The user presents a username, server and password. Users must present both the VPN username and the e-mail username, separated by the VPN Name Delimiter, only if the usernames are different from each other.

- Certificate—Certificate authentication does not work for e-mail proxies in the current security appliance software release.
- Piggyback HTTPS—Select to require piggyback authentication.

  This authentication scheme requires a user to have already established a WebVPN session. The user presents an e-mail username only. No password is required. Users must present both the VPN username and the e-mail username, separated by the VPN Name Delimiter, only if the usernames are different from each other.

  SMTPS e-mail most often uses piggyback authentication because most SMTP servers do not allow users to log in.

Note IMAP generates a number of sessions that are not limited by the simultaneous user count but do count against the number of simultaneous logins allowed for a username. If the number of IMAP sessions exceeds this maximum and the WebVPN connection expires, a user cannot subsequently establish a new connection. There are several solutions:

- The user can close the IMAP application to clear the sessions with the security appliance, and then establish a new WebVPN connection.
- The administrator can increase the simultaneous logins for IMAP users (Configuration > Features > VPN > General > Group Policy > Edit Group Policy > General).
- Disable HTTPS/Piggyback authentication for e-mail proxy.

- Mailhost—(SMTPS only) Select to require mailhost authentication. This option appears for SMTPS only because POP3S and IMAP4S always perform mailhost authentication. It requires the user's e-mail username, server and password.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Default Servers

Configuration > Features > VPN > E-mail Proxy > Default Servers

This panel lets you identify proxy servers to the security appliance. Enter the IP address and port of the appropriate proxy server.

Fields

- POP3S/IMAP4S/SMTPS Default Server—Let you configure a default server, port and non-authenticated session limit for e-mail proxies.
- Name or IP Address—Type the DNS name or IP address for the default e-mail proxy server.

- Port—Type the port number on which the security appliance listens for e-mail proxy traffic. Connections are automatically allowed to the configured port. The e-mail proxy allows only SSL connections on this port. After the SSL tunnel establishes, the e-mail proxy starts, and then authentication occurs.

  For POP3s the default port is 995, for IMAP4S it is 993, and for SMTPS it is 988.

- Enable non-authenticated session limit—Select to restrict the number of non-authenticated e-mail proxy sessions.

  E-mail proxy connections have three states:

  1. A new e-mail connection enters the “unauthenticated” state.
  2. When the connection presents a username, it enters the “authenticating” state.
  3. When the security appliance authenticates the connection, it enters the “authenticated” state.

  This feature lets you set a limit for sessions in the process of authenticating, thereby preventing DOS attacks. When a new session exceeds the set limit, the security appliance terminates the oldest non-authenticating connection. If there are no non-authenticating connections, the oldest authenticating connection is terminated. The does not terminate authenticated sessions.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Delimiters

Configuration > Features > VPN > E-mail Proxy > Delimiters

This panel lets you configure username/password delimiters and server delimiters for e-mail proxy authentication.

Fields

- POP3S/IMAP4S/SMTPS Delimiters—Let you configure username/password and server delimiters for each of the e-mail proxies.
  - Username/Password Delimiter—Select a delimiter to separate the VPN username from the e-mail username. Users need both usernames when using AAA authentication for e-mail proxy and the VPN username and e-mail username are different. Users enter both usernames, separated by the delimiter you configure here, and also the e-mail server name, when they log in to an e-mail proxy session.

Note Passwords for WebVPN e-mail proxy users cannot contain characters that are used as delimiters.

– Server Delimiter—Select a delimiter to separate the username from the name of the e-mail server. It must be different from the VPN Name Delimiter. Users enter both their username and server in the username field when they log in to an e-mail proxy session.

  For example, using : as the VPN Name Delimiter and @ as the Server Delimiter, when logging in to an e-mail program via e-mail proxy, the user would enter their username in the following format: vpn_username:e-mail_username@server.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

CHAPTER 32

Configuring SSL Settings

SSL

Configuration > Properties > SSL

The security appliance uses the Secure Sockets Layer (SSL) protocol and its successor, Transport Layer Security (TLS) to achieve secure message transmission for both ASDM and WebVPN sessions. The SSL window lets you configure SSL versions for clients and servers and encryption algorithms. It also lets you apply previously configured trustpoints to specific interfaces, and to configure a fallback trustpoint for interfaces that do not have an associated trustpoint.

Fields

- Server SSL Version—Choose to specify the SSL/TLS protocol version the security appliance uses when acting as a server. You can make only one selection.

Options for Server SSL versions include the following:

Any	The security appliance accepts SSL version 2 client hellos, and negotiates either SSL version 3 or TLS version 1.
Negotiate SSL V3	The security appliance accepts SSL version 2 client hellos, and negotiates to SSL version 3.
Negotiate TLS V1	The security appliance accepts SSL version 2 client hellos, and negotiates to TLS version 1.
SSL V3 Only	The security appliance accepts only SSL version 3 client hellos, and uses only SSL version 3.
TLS V1 Only	The security appliance accepts only TLSv1 client hellos, and uses only TLS version 1.

Note To use WebVPN port forwarding, you must select Any or Negotiate SSL V3. The issue is that JAVA only negotiates SSLv3 in the client Hello packet when you launch the Port Forwarding application.

- Client SSL Version—Choose to specify the SSL/TLS protocol version the security appliance uses when acting as a server. You can make only one selection.

Options for Client SSL versions include the following:

any	The security appliance sends SSL version3 hellos, and negotiates either SSL version 3 or TLS version 1.
sslv3-only	The security appliance sends SSL version 3 hellos, and accepts only SSL version 3.
tlsv1-only	The security appliance sends TLSv1 client hellos, and accepts only TLS version 1.

- Encryption—Lets you set SSL encryption algorithms.
  - Available Algorithms—Lists the encryption algorithms the security appliance supports that are not in use for SSL connections. To use, or make active, an available algorithm, highlight the algorithm and click **Add**.
  - Active Algorithms—Lists the encryption algorithms the security appliance supports and is currently using for SSL connections. To discontinue using, or change an active algorithm to available status, highlight the algorithm and click **Remove**.
  - Add/Remove—Click to change the status of encryption algorithms in either the Available or Active Algorithms columns.
  - Move Up/Move Down—Highlight an algorithm and click these buttons to change its priority. The security appliance attempts to use an algorithm
- Trustpoints**—**Lets you select a fallback trustpoint, and displays configured interfaces and the configured trustpoints associated with them. To enroll a trustpoint, go to **Configuration > Properties > Certificate > Enrollment**.
  - Fallback Trustpoint—Click to select a trustpoint to use for interfaces that have no trustpoint associated with them. If you select None, the security appliance uses the default RSA key-pair and certificate.

Note The trustpoint must have a certificate associated with it to display in this drop-down list.

  - Interface and Trustpoint columns—Display configured interfaces and the trustpoint, if any, for the interface.
  - Edit—Click to change the trustpoint for the highlighted interface.
- Apply**—**Click to apply your changes.
- Reset—Click to remove changes you have made and reset SSL parameters to the values that they held when you opened the window.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Edit SSL Trustpoint

Configuration > Properties > SSL > Edit SSL Trustpoint

Fields

- Interface—Displays the name of the interface you are editing.
- Enrolled Trustpoint—Click to select a previously enrolled trustpoint to associate with the named interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

CHAPTER 33

Configuring Certificates

Digital certificates provide digital identification for authentication. A digital certificate contains information that identifies a device or user, such as the name, serial number, company, department, or IP address. CAs issue digital certificates in the context of a PKI, which uses public-key/private-key encryption to ensure security. CAs are trusted authorities that "sign" certificates to verify their authenticity, thus guaranteeing the identity of the device or user.

A *CA certificate* is one used to sign other certificates. A CA certificate that is self-signed is called a *root certificate*; one issued by another CA certificate is called a *subordinate certificate*. CAs also issue *identity certificates*, which are the certificates for specific systems or hosts.

For authentication using digital certificates, there must be at least one identity certificate and its issuing CA certificate on a security appliance, which allows for multiple identities, roots and certificate hierarchies.

For More Information

Authenticating, Enrolling for, and Managing Digital Certificates

Authentication

Configuration > Properties > Certificate > Authentication

The Authentication panel lets you authenticate a CA certificate, which associates the CA certificate with a trustpoint and installs it on the security appliance. You can edit an existing trustpoint configuration or you can create a new one.

If the trustpoint you select is configured for manual enrollment, you should obtain the CA certificate manually and import it here. If the trustpoint you select is configured for automatic enrollment, the security appliance uses the SCEP protocol to contact the CA, and then automatically obtains and installs the certificate.

Fields

- **Trustpoint Name**—Displays a list containing the trustpoints available from which to obtain the CA certificate. Click a trustpoint in the list and edit its configuration, or add a new trustpoint.
- **Edit**—Click to modify a trustpoint configuration currently appearing in the Trustpoint Name box.
- **New**—Add a new trustpoint configuration to the list.

- **Fingerprint**—Specify a key consisting of alphanumeric characters the security appliance uses to authenticate the CA certificate. If you provide a fingerprint, the security appliance compares it to the computed fingerprint of the CA certificate and accepts the certificate only if the two values match. If there is no fingerprint, the security appliance accepts the certificate without one.
- **Import from a file**—For manual enrollment only, identify a file from which to import the certificate. You can type the pathname of the file in the box or you can click Browse and search for the file.
  - Browse—Displays the Load Certificate File dialog box that lets you navigate to the file containing the certificate.
- **Enter the certificate text in base64 format**—For manual enrollment, enter the trustpoint configuration in base64 format.
- **Authenticate**—Complete the authentication procedure.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

For More Information

Authenticating, Enrolling for, and Managing Digital Certificates

Enrollment

Configuration > Properties > Certificate > Enrollment

The Enrollment panel lets you select a trustpoint configuration from the list, edit a trustpoint configuration or create a new one. However, for automatic enrollment, you cannot generate an enrollment request until you have authenticated the CA certificate.

For automatic enrollment, the security appliance contacts the CA using SCEP protocol, obtains the identity certificates, and installs them on the device. For manual enrollment, an enrollment request dialog box appears containing the certificate enrollment request. Use this enrollment request to obtain the identity certificate from the management interface of the CA. The identity certificate obtained must be in base64 or hexadecimal format. You can then import it in the Import Certificate dialog box.

Fields

- **Trustpoint Name**—Specify the trustpoint for which to generate the enrollment request. Select the name from a list, edit the name currently appearing in the box, or add a new trustpoint configuration.
- **Edit**—Modify the trustpoint configuration currently appearing in the Trustpoint Name box.
- **New**—Add a new trustpoint configuration to the list.
- **Enroll**—Initiate the enrollment process with the CA.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

For More Information

Authenticating, Enrolling for, and Managing Digital Certificates

Import Certificate

Configuration > Properties > Certificate > Import Certificate

The Import Certificate panel lets you install the device certificate that you received from the CA during manual enrollment. To import certificates from a CA, there should be a CA certificate associated with the selected trustpoint. If not, the security appliance displays a warning.

Fields

- **Trustpoint Name**—Specify the name of the trustpoint from which you received the certificate. Select the name from a list, edit the name currently appearing in the box, or add a new trustpoint configuration.
- **Edit**—Modify the trustpoint configuration currently appearing in the Trustpoint Name box.
- **New**—Add a new trustpoint configuration to the list.
- **Import from a file**—Identify a file from which to import the identity certificate. You can type the pathname of the file in the box or you can click Browse and search for the file.
  - **Browse**—Displays the Load CA certificate file dialog box that lets you navigate to the file containing the certificate.
- **Enter the certificate text in base64 format**—For manual enrollment, lets you use cut and paste to transfer the certificate data to this security appliance from the source exported.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Key Pair

Configuration > Properties > Certificate > Key Pair

RSA key pairs are required to enroll for identity certificates. The security appliance supports multiple key pairs.

Fields

- **Key-pair Name**—Displays the name given to the key pair(s).
- **Type**—Displays the type, which is RSA.
- **Usage**—Displays how an RSA key pair is to be used. There are two types of usage for RSA keys: general purpose, the default, and special. When you select Special, the security appliance generates two key pairs, one for signature use and one for encryption use. This implies that two certificates for the corresponding identity are required.
- **Size**—Displays the modulus size of the key pair(s): 512, 768, 1024, and 2048. The default modulus size is 1024.
- **Add**—Opens the Add Key Pair dialog box.
- **Show Details**—Displays the name, date generated, type, modulus size, usage and DER-encoded key data.
- **Delete**—Deletes the selected key pair.
- **Refresh**—Updates the display.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Add Key Pair

Configuration > Properties > Certificate > Key Pair > Add Key Pair

The Add Key Pair dialog box lets you add a new key pair to the list of key pairs.

Fields

- **Name**—Specify a name for the key pair(s): the default key <Default-RSA-Key> or a specific key. The security appliance uses the default key pair when a trustpoint has no key pairs configured.
- **Size**—Specify the modulus size of the key pair(s): 512, 768, 1024, and 2048. The default modulus size is 1024.
- **Type**—Specify the type, which can be RSA only.
- **Usage**—Specify how the key pair is to be used. There are two types of usage for RSA keys: general purpose, the default, or special. When you click Special, the security appliance generates two key pairs, one for signature use and one for encryption use. This implies that two certificates for the corresponding identity are required.
- **Generate Now**—Generate the key pair.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Key Pair Details

Configuration > Properties > Certificate > Key Pair > Show Details

The Key-pair Details dialog box displays information about the selected key-pair.

Fields

- **Key Pair**—Displays the name given to the key pair.
- **Generation Time**—Displays time and date that the key was generated.
- **Type**—Displays the type of key pair (RSA).
- **Size**—Displays the modulus size. For RSA keys, the size can be 512, 768, 1024, or 2048. The default modulus size is 1024.
- **Usage**—Displays how an RSA key pair is to be used. There are two types of usage for RSA keys: general purpose, the default, and special. When the purpose of the key pair is Special, the security appliance generates two key pairs, one for signature use and one for encryption use. This implies that two certificates for the corresponding identity are required.
- **Key Data**—Displays the DER-encoded key data.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Manage Certificate

Configuration > Properties > Certificate > Manage Certificates

The Manage Certificates panel displays all of your certificates in a table and lets you add/edit a certificate, display certificate information, refresh a display and delete certificates from the security appliance.

Fields

- **Subject**—Identifies the owner of the certificate.

- **Type**—Identifies the type: CA, RA general, RA encryption, RA signature, identity.
- **Trustpoint**—Identifies the trustpoint.
- **Status**—Identifies the status: Available or Pending:
  - **Available** means that the CA has accepted the enrollment request and has issued an identity certificate.
  - **Pending** means that the enrollment request is still in process and that the CA has not issued the identity certificate yet.
- **Usage**—Identifies how the certificate is used: signature, general purpose, or encryption.
- **Add**—Displays the Add Certificate dialog box, which lets you add CA/RA/Identity certificates onto the security appliance. You can use this dialog box to import a certificate from a file you have exported or use cut and paste to enter a certificate onto the security appliance.
- **Show Details**—Displays the Certificate Details dialog box, which shows the following information about the selected certificate:
  - **General**—Displays the values for type, serial number, status, usage, CRL distribution point, and the time within which the certificate is valid. This applies to both available and pending status.
  - **Subject**— Displays the X.500 fields of the subject DN or certificate owner and their values. This applies only to available status.
  - **Issuer**—Displays the X.500 fields of the entity that granted the certificate. This applies only to available status.
- **Refresh**—Renews the display of the table in the Manage Certificates panel.
- **Delete**—Displays the Delete Certificate dialog box that asks you to confirm the certificate removal. If you delete a CA certificate, the security appliance deletes all the associated identity certificates as well.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

For More Information

Authenticating, Enrolling for, and Managing Digital Certificates

Add Certificate

Configuration > Properties > Certificate > Manage Certificates > Add Certificate

The Add Certificate dialog box lets you manually add CA/RA/Identity certificates.

Fields

- **Trustpoint Name**—Specify the certificate to add to the Manage Certificates table.

- **Edit**—Modify the trustpoint configuration currently appearing in the Trustpoint Name box.
- **New**—Add a new trustpoint configuration to the list.
- **Certificate Type**—Specify the type: CA, RA general, RA encryption, RA signature, Identity.
- **Serial Number**—Include the serial number of the security appliance in the certificate.
- **Import from a file**—Identify a file from which to import the certificate. You can type the pathname of the file in the box or you can click Browse and search for the file.
  - **Browse**—Display the Add Certificate dialog box that lets you navigate to the file containing the certificate.
- **Enter the certificate text in base64 format**—Lets you use cut and paste to transfer the certificate data to this security appliance from the source text that was exported, which should be in hexadecimal format only.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Trustpoint

A trustpoint represents a CA/identity pair and contains the identity of the CA, CA-specific configuration parameters, and an association with one enrolled identity certificate.

Configuration

Configuration > Properties > Certificate > Trustpoint > Configuration

The Configuration panel lets you identify a CA, which can be a root CA, and have a self-signed certificate that contains its own public key. In the Configuration panel, you can add, edit, or delete a CA as a trustpoint, and request a CRL.

Fields

- **Trustpoint Name**—Displays the name of the trustpoint, for example, an IP address or a hostname.
- **Device Certificate Subject**—Displays the subject DN owning the certificate for the security appliance system.
- **CA Certificate Subject**—Displays the subject name of the CA certificate.
- **Add**—Opens the Add Trustpoint Configuration dialog box.
- **Edit**—Opens the Edit Trustpoint Configuration dialog box.
- **Delete**—Removes the selected trustpoint.
- **Request CRL**—Retrieves the Certificate Revocation List for the selected trustpoint. To view it, see Monitoring > Properties > CRL.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Add/Edit Trustpoint Configuration > Enrollment Settings Tab

Configuration > Properties > Certificate > Trustpoint > Configuration > Add/Edit Trustpoint Configuration > Enrollment Settings Tab

The **Enrollment Settings** tab lets you add a trustpoint to the trustpoint table. The **Edit Trustpoint Configuration > Enrollment Settings** tab lets you modify information about the selected trustpoint.

Fields

- **Trustpoint Name—**Specify the name of the trustpoint corresponding to a CA. For example, this can be an IP address or a hostname.
- **Generate a self-signed certificate on enrollment**—Click to generate a self-signed device certificate for the security appliance during enrollment. This provides a way to create self-signed certificates for use when terminating SSL connections. This feature is not checked by default. When this option is checked, you can configure only the key pair and the certificate parameters.
- **Key Pair**—Select a previously defined key pair in the list. Before you add a trustpoint, you should configure a key pair. So if this list is empty, you can add the key pair by selecting **New Key Pair**.
- **Show Details**—Display information about the key pair including its name, when it was generated, its type (RSA), its modulus, its usage (general purpose or special) and the key data in DER-encoded format.
- **New Key Pair**—Open the **Add Key Pair** dialog box, which lets you enter a name, size, type, and usage for a new key pair.
- **Challenge Password**—Specify a challenge phrase that is registered with the CA during enrollment.
- **Confirm Challenge Password**—Verify the challenge password.
- **Use manual enrollment**—Specify intention to generate a PKCS10 certification request. The CA issues a certificate to the security appliance based on the request and the certificate is installed on the security appliance by importing the new certificate.
- **Use automatic enrollment**—Specify intention to use SCEP mode. When the indicated trustpoint is configured for SCEP enrollment, the security appliance then downloads the certificates using the SCEP protocol.
- **Enrollment URL**—Specify the name of the URL for automatic enrollment. The maximum length is 1000 characters (effectively unbounded).
- **Retry Period**—After requesting a certificate, the security appliance waits to receive a certificate from the CA. If the security appliance does not receive a certificate within the specified retry period, it sends another certificate request. Use this field to specify the number of minutes between attempts to send an enrollment request; the valid range is 1- 60 minutes. The default value is 1.

- **Retry Count**—After requesting a certificate, the security appliance waits to receive a certificate from the CA. If the security appliance does not receive a certificate within the specified retry period, it sends another certificate request. The security appliance repeats the request until either it receives a response or reaches the retry count specified. Use this field to specify the maximum number of attempts to send an enrollment request, the valid range is 0, 1-100 retries. The default value is 0, which means an unlimited number of retries.
- **Certificate Parameters**—Display the **Certificate Parameters** dialog box, which lets you specify attributes and their values to include in the certificate during enrollment, such as subject DN, FQDN, and so on.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Add/Edit Key Pair

Configuration > Properties > Certificate > Trustpoint > Configuration > Add/Edit Trustpoint Configuration > Enrollment Settings > Add/Edit Key Pair

The **Add Key Pair** dialog box lets you add a new key pair to the list of key pairs.

Fields

- **Name**—Specify a name for the key pair(s): the default key <Default-RSA-Key> or a specific key. The security appliance uses the default key pair when a trustpoint has no key pairs configured.
- **Size**—Specify the modulus size of the key pair(s): 512, 768, 1024, and 2048. The default modulus size is 1024.
- **Usage**—Specify how the key pair is to be used. There are two types of usage for RSA keys: general purpose, the default, or special. When you click **Special**, the security appliance generates two key pairs, one for signature use and one for encryption use. This implies that two certificates for the corresponding identity are required.
- **Generate Now**—Generate the key pair.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Certificate Parameters

Configuration > Properties > Certificate > Trustpoint > Configuration > Add/Edit Trustpoint Configuration > Enrollment Settings > Certificate Parameters

The **Certificate Parameters** dialog box lets you specify the subject DN, FQDN, IP address to include during enrollment. Use this dialog box to include the device serial number.

Fields

- **Subject DN—**Specify the attributes and values to use for the X.500 name of the subject. The subject is the owner of the certificate.
  - Click **Edit** to display the **Edit DN** dialog box to select the attributes and values for the **Subject DN**.
- **FQDN**—Include the fully qualified domain name in the Subject Alternative Name extension of the certificate. The FQDN is the part of a URL that completely identifies the server program that a request is addressed to; for example www.examplesite.com.
- **E-mail**—Include the indicated e-mail address in the Subject Alternative Name extension of the certificate.
- **IP Address**—Include the indicated IP address in the Subject Alternative Name extension of the certificate.
- **Include device serial number**—Include the security appliance's serial number in the certificate during enrollment.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Edit DN

Configuration > Properties > Certificate > Trustpoint > Configuration > Add/Edit Trustpoint Configuration > Enrollment Settings > Certificate Parameters > Edit DN

Edit DN

Select one of the following attributes in the **Attributes** list, type the value in the **Value** box, and click **Add**. Select as many as are needed.

Fields

- **Common Name (CN)**—An individual's given name; for example, Pat.
- **Department (OU)**—Organizational Unit or a subgroup of a larger organization such as an enterprise or a university; for example, Geology department.
- **Company Name (O)**—Organization such as an enterprise or university; for example, University of Oz.
- **Country (C)**—Two-letter designation for a specific country; for example, OZ.

- **State (St)**—State or Province within a country; for example, Kansas.
- **Location (L)**—Address of the subject; for example, 49 Wizard St.
- **Email Address (EA)**—Pat@univoz.org.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Add/Edit Trustpoint Configuration > Revocation Check Tab

Configuration > Properties > Certificate > Trustpoint > Configuration > Add/Edit Trustpoint Configuration > Revocation Check Tab

The **Revocation Check** tab lets you specify whether to check certificates for revocation, and if you do, the method to use.

When a certificate is issued, it is valid for a fixed period of time. Sometimes a CA revokes a certificate before this time period expires; for example, due to security concerns or a change of name or association. CAs periodically issue a signed list of revoked certificates. Enabling revocation checking forces the security appliance to check that the CA has not revoked a certificate every time it uses that certificate for authentication.

The security appliance supports two methods for checking revocation status: CRL and OCSP.

Fields

- **Do not check certificates for revocation**—Select if you want the security appliance not to check certificates for revocation status.
- **Check certificates for revocation**—Select to have the security appliance check certificates for revocation status. You must also specify at least one revocation method.
- **Revocation methods**—Specify the revocation methods to use in checking certificates. If you specify more than one method, the security appliance applies methods in the order you set here. It uses the second method only if the first returns an error, for example, if the server is unavailable. Methods available include CRL and OCSP.
  - CRL—The security appliance retrieves, parses, and caches the complete certificate revocation list to determine the status of a certificate.
  - OCSP— The security appliance localizes certificate status on a Validation Authority, which it can query for the status of a specific certificate.
- **Add**—Click either CRL or OCSP in the left to add it as a revocation checking method.
- **Remove**—Click either CRL or OCSP in the right to remove it as a revocation checking method.
- **Move Up/Move Down**—Use these buttons to have the security appliance first use the method you prefer.
- **Consider certificate valid if revocation checking returns errors**—Check to have the security appliance accept a certificate even if errors occur during a revocation check.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Add/Edit Trustpoint Configuration > CRL Retrieval Policy Tab

Configuration > Properties > Certificate > Trustpoint > Configuration > Add/Edit Trustpoint Configuration > CRL Retrieval Policy Tab

The **CRL Retrieval Policy** tab lets you specify whether to retrieve CRLs from CRL DPs or from URLs listed in the **Static URLs** table.

Fields

- **Use CRL Distribution Point from the certificate**—Click to retrieve CRLs from the distribution point listed in the certificate.
- **Use Static URLs configured below**—Click to add up to five URLs from which the security appliance should attempt to obtain a CRL.
- **Add**—Displays the **Add Static URL** box. Use this box to add up to five URLs.
  - **URL:**—Select URL type: HTTP, LDAP, or SCEP.
  - **://**—Type the location that distributes the CRLs.
- **Edit**—Display the **Edit Static URL** box for you to modify the selected URL.
- **Delete**—Remove the selected URL.
- **Move Up**—Move the selected URL up in the table, until it is at the top.
- **Move Down**—Move the selected URL down in the table, until it is at the bottom.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Add/Edit Static URL

Configuration > Properties > Certificate > Trustpoint > Configuration > Add/Edit Trustpoint Configuration > CRL Retrieval Policy Tab > Add/Edit Static URL

Fields

- **URL:**—Select URL type: HTTP, LDAP, or SCEP.

– **://**—Type the location that distributes the CRLs.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Add/Edit Trustpoint Configuration > CRL Retrieval Method Tab

Configuration > Properties > Certificate > Trustpoint > Configuration > Add/Edit Trustpoint Configuration > CRL Retrieval Method Tab

The **CRL Retrieval Method** tab lets you specify how to retrieve CRLs, including LDAP, HTTP and SCEP. You can enable all methods. If you enable several methods, ASDM uses them in the order you specify.

Fields

- **Enable Lightweight Directory Access Protocol (LDAP)**—Check to enable.

  Specify LDAP parameters as follows:

  – **Name**—Identify the person who has access to the CRL on the server.
  – **Password**—Specify a password for the person listed under **Name**.
  – **Confirm Password**—Verify the password.
  – **Default Server**—Specify the hostname or IP address of the LDAP server.
  – **Default Port**—Specify the LDAP server port number. The default is 389.
- **Enable HTTP**—Specify HTTP as a protocol to use for CRL retrieval.
- **Enable Simple Certificate Enrollment Protocol (SCEP)**—Use the same method of retrieving the CRL as for enrollment, but not at enrollment time.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Add/Edit Trustpoint Configuration > OCSP Rules Tab

Configuration > Properties > Certificate > Trustpoint > Configuration > Add/Edit Trustpoint Configuration > OCSP Rules Tab

The OCSP Rules tab lets you configure OCSP certificate matching rules. These rules provide flexibility in that you can assign OCSP server URLs via a trustpoint

While you can configure multiple match rules for a trustpoint, only one match rule within a trustpoint can apply to a certificate map.

Fields

- **Certificate Map**—Displays the name of the certificate map to match to this OCSP rule. Certificate maps match user permissions to specific fields in a certificate. You must configure the certificate map before you configure OCSP rules (Configuration > VPN > IKE > Certificate Group Matching > Rules).
- **Trustpoint**—Displays the name of the trustpoint the security appliance uses to validate responder certificates.
- **Index**—Displays the priority number for the rule. The security appliance examines OCSP rules in priority order, and applies the first one that matches.
- **URL**—Specifies the URL for the OCSP server for this trustpoint.
- **Add**—Click to add a new OCSP rule.
- **Edit**—Click to edit an existing OCSP rule.
- **Delete**—Click to delete an OCSP rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Add/Edit Trustpoint OCSP Rule dialog box

Configuration > Properties > Certificate > Trustpoint > Configuration > Add/Edit Trustpoint Configuration > OCSP Rules Tab > Add/Edit Trustpoint OCSP Rule

You can configure OCSP rules for a trustpoint that override the OCSP server URL specified within the AuthorityInfoAccess (AIA) field of the remote user certificate.

- **Certificate Map**—Select the name of the certificate map to match to this OCSP rule. Certificate maps match user permission groups to specific fields in a certificate. Their function with OCSP is to let the security appliance access a particular OCSP server for revocation status, as well as to let you specify a trustpoint to validate the responder certificate. This lets you check revocation status via a trustpoint other than the trustpoint authenticating the remote user certificate.

  You must configure the certificate map before you configure OCSP rules (Configuration > VPN > IKE > Certificate Group Matching > Rules).
- **Trustpoint**—Select the trustpoint that you want to use for this OCSP rule. You must have already configured this trustpoint.
- **Index**—Enter a number to determine the execution order of the match rules. The security appliance searches the match rules lowest to highest, according to this index, and applies the first rule that matches.

- **URL**—Specify the URL for the OCSP server for this trustpoint.

  The security appliance uses OCSP servers in this order:

  1. OCSP URL in a match certificate override rule (as configured here)

  2. OCSP URL configured in the Add/Edit Trustpoint Configuration > Advanced Tab > OCSP Options attribute

  3. AIA field of remote user certificate

  It you do not set this URL attribute, the OCSP server specified the Advanced Tab > OCSP Options attribute applies, and if that is not set, the OCSP server in the Authority Info Access (AIA) extension of the remote user certificate applies. If the AIA does not have an AIA extension and you do not set a valid OCSP server here or in the Advanced tab, revocation status checking fails.

  The security appliance supports only HTTP URLs, and you can specify only one URL per trustpoint.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Add/Edit Trustpoint Configuration > Advanced Tab

Configuration > Properties > Certificate > Trustpoint > Configuration > Add/Edit Trustpoint Configuration > Advanced Tab

The **Advanced** tab lets you specify CRL and OCSP options. When a certificate is issued, it is valid for a fixed period of time. Sometimes a CA revokes a certificate before this time period expires; for example, due to security concerns or a change of name or association. CAs periodically issue a signed list of revoked certificates. Enabling revocation checking forces the security appliance to check that the CA has not revoked the certificate being verified.

The security appliance supports two methods of checking revocation status: CRL and OCSP.

Fields

- **CRL Options**
  - **Cache Refresh Time**—Specify the number of minutes between cache refreshes. The default number of minutes is 60. The range is 1-1440.

    To avoid having to retrieve the same CRL from a CA repeatedly, The security appliance can store retrieved CRLs locally, which is called CRL caching. The CRL cache capacity varies by platform and is cumulative across all contexts. If an attempt to cache a newly retrieved CRL would exceed its storage limits, the security appliance removes the least recently used CRL until more space becomes available.
  - **Enforce next CRL update**—Require valid CRLs to have a Next Update value that has not expired. Clearing the box allows valid CRLs with no Next Update value or a Next Update value that has expired.
- **OCSP Options**

  - **Server URL:**—Enter the URL for the OCSP server. The security appliance uses OCSP servers tin the following order:

    1. OCSP URL in a match certificate override rule (Add/Edit Trustpoint Configuration > OCSP Rules tab)

    2. OCSP URL configured in this OCSP Options attribute

    3. AIA field of remote user certificate

  - **Disable nonce extension**—By default the OCSP request includes the nonce extension, which cryptographically binds requests with responses to avoid replay attacks. It works by matching the extension in the request to that in the response, ensuring that they are the same. Disable the nonce extension if the OCSP server you are using sends pre-generated responses that do not contain this matching nonce extension.

- **Accept certificates issued by this trustpoint**—Specify whether or not the security appliance should accept certificates from **Trustpoint Name**.
- **Accept certificates issued by the subordinate CAs of this trustpoint**
- **Use the configuration of this trustpoint to validate any remote user certificate issued by the CA corresponding to this trustpoint**—When enabled, the configuration settings active when a remote user certificate is being validated can be taken from this trustpoint if this trustpoint is authenticated to the CA that issued the remote certificate.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Export

Configuration > Properties > Certificate > Trustpoint > Export

The **Export** panel lets you export a trustpoint configuration with all associated keys and certificates in PKCS12 format, which must be in base64 format. An entire trustpoint configuration includes the entire chain (root CA certificate, identity certificate, key pair) but not enrollment settings (subject name, FQDN and so on). This feature is commonly used in a failover or load balancing configuration to replicate trustpoints across a group of security appliances; for example, remote access clients calling in to a central organization that has several units to service the calls. These units must have equivalent trustpoint configurations. In this case, an administrator can export a trustpoint configuration and then import it across the group of security appliances.

Fields

- **Trustpoint Name—**Click a trustpoint in the list and edit its configuration, or add a new trustpoint configuration.
- **Edit**—Modify the trustpoint configuration currently appearing in the **Trustpoint Name** box
- **New**—Add a new trustpoint configuration to the list.
- **Encryption Passphrase**—Specify the passphrase used to encrypt the PKCS12 file for export.

- **Confirm Passphrase**—Verify the encryption passphrase.
- **Export to a file**—Specify the name of the PKCS12-format file to use in exporting the trustpoint configuration; PKCS12 is the public key cryptography standard, which can be base64 encoded or hexadecimal.
  - **Browse**—Display the **Select a File** dialog box that lets you navigate to the file to which you want to export the trustpoint configuration.
- **Display the trustpoint configuration in PKCS12 format—**Display the **Export Trustpoint Configuration** dialog box, which displays the trustpoint configuration in a text box. You can use cut and paste to extract the data and place it in the window of the **Import** panel. To exit, click **OK**.
- **Export**—Export the trustpoint configuration.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Import

Configuration > Properties > Certificate > Trustpoint > Import

The **Import** panel lets you install an entire trustpoint configuration in PKCS12 format. An entire trustpoint configuration includes the entire chain (root CA certificate, RA certificate, identity certificate, key pair) but not enrollment sets (subject name, FQDN and so on). This feature is commonly used in a failover or load balancing configuration to replicate trustpoints across a group of security appliances; for example, remote access clients calling in to a central organization that has several units to service the calls. These units must have equivalent trustpoint configurations. In this case, an administrator can export a trustpoint configuration and then import it across the group of security appliances.

Fields

- **Trustpoint Name—**Identify the trustpoint. When importing from another security appliance for failover or load balancing, you can use the same trustpoint name as the security appliance from which the trustpoint configuration was exported. However make sure that a trustpoint/key pair with the same name does not already exist.
- **Decryption Passphrase—**Specify the encryption passphrase specified during the export of the trustpoint configuration.
- **Confirm Passphrase—**Verify the passphrase.
- **Import from a file—**Identify a file from which to import the certificate. The text imported from a file should be PKCS12 data, in either base64 or hexadecimal format. You can type the pathname of the file in the box or you can click **Browse** and search for the file.
  - **Browse**—Display the **Load Certificate File** dialog box that lets you navigate to the file containing the trustpoint configuration.

- **Enter the trustpoint configuration in PKCS12 format—**lets you paste the trustpoint configuration in PKCS12 format, which can be in either base64 or hexadecimal format. In this case, you use cut and paste to enter the data into the text box.
- **Import**—Import the trustpoint configuration.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	•

Authenticating, Enrolling for, and Managing Digital Certificates

This section describes how to enroll for a digital certificate. Once enrolled, you can use the certificate for authenticating the device to VPN and SSL peers.

Summary of Configuration Steps

Here are the basic steps for enrolling with a CA and getting an identity certificate to use for authenticating tunnels. This example shows both automatic (SCEP) enrollment and manual enrollment. For information on fields not defined in this procedure, click the **Help** button.

1. Generating a key pair for the identity certificate. The key pair is RSA.
2. Creating a trustpoint.
3. Configuring an enrollment URL.
4. Authenticating the CA.
5. Enrolling with the CA, which places an identity certificate onto the security appliance.

Note Authenticating and Enrolling are two separate phases of the process. You must authenticate. Then you can enroll using either automatic enrollment or manual enrollment.

Generating the Key Pair

Begin by generating a key pair for the certificate. Generated key pairs are identified by labels that you provide when you configure the key pair. RSA Key pairs come in two types: general purpose and usage. General purpose is the default type and generates a single pair of keys. Usage type generates two key pairs, one for signature use and one for encryption use, thus requiring two certificates for the corresponding identity.

To generate an RSA key pair using ASDM, follow this procedure:

Step 1 Under **Configuration > Features > Device Administration > Certificate > Key Pair**, click **Add**.

Step 2 Configure the information in the **Add Key Pair** dialog box:

Step 3 Click **Generate Now**.

Step 4 To view the key pair generated, click **Show Details**. ASDM displays information about the key pair.

Enrolling for a Certificate Using Automatic Enrollment (SCEP)

Create a trustpoint. A trustpoint represents a CA/identity pair and contains the identity of the CA, specific configuration parameters, and an association with one enrolled identity certificate.

To create a trustpoint, follow these steps:

Step 1 Under **Configuration > Features > Device Administration > Certificate >Trustpoint > Configuration**, click **Add**.

Step 2 Configure the basic information in the **Add Trustpoint Configuration** dialog box. For all other parameters, you can accept the default values.

a. **Trustpoint Name**—Type the trustpoint name in the **Trustpoint Name** box.

b. **Enrollment URL**—In the **Enrollment Settings** panel, under the **Enrollment Mode** group box, for SCEP enrollment, click **Use automatic enrollment**. Then type the enrollment URL in the box. For example, type 10.20.30.40/cgi-bin/pkiclient.exe.

c. If you want password verification for the certificate, type the password into the **Challenge Password** and **Confirm Password** boxes. If you need to revoke the certificate, you can provide this password to the CA administrator to identify that you are the certificate owner. This password is not saved in the configuration, so you should make a note of it.

Step 3 Configure the configuration parameters next. At the very least, you need to configure a subject name for the certificate using X.500 fields; for example, common name (CN) and organizational unit (OU).

a. In the **Enrollment Settings** panel, select the key pair you configured for this trustpoint in the **Key Pair** list.

b. In the **Enrollment Settings** panel, click **Certificate Parameters**.

c. To add subject DN values, click **Edit** in the **Certificate Parameters** dialog box.

d. In the **Edit DN** box under **DN Attribute to be Added**, select an attribute in the **Attribute** list and type a value in the **Value** box. Then click **Add**. For example, first select **Command Name (CN)** and type Pat in the **Value** box; then select **Department (OU)** and type Engineering in the **Value** box.

e. After entering all subject DN information, click **OK.**

f. Optionally type values for **FQDN**, **E-mail**, and **IP Address**, and check the **Include device serial number** option.

g. Click **OK**.

Step 4 Click **Apply**. If you have preview commands checked, ASDM displays the CLI commands based on the ASDM configuration for you to either send or cancel. Click **Send**. Do this for all features you configure using this procedure.

Authenticating to the CA

Authenticating to the CA puts the CA certificate onto the security appliance. If you configure the trustpoint for SCEP enrollment, the CA certificate is downloaded through SCEP. If not, you must paste the CA certificate into the text box or point to the file with the browse button. This section shows SCEP enrollment.

To authenticate to the CA, follow these steps:

Step 1 Under **Configuration > Features > Device Administration > Certificate > Authentication**, select the name of the trustpoint in the **Trustpoint Name** list.

Step 2 Click **Authenticate**.

Step 3 When ASDM displays the **Authentication Successful** dialog, click **OK**.

Enrolling with the CA

After you have configured the trustpoint and authenticated with it, you can enroll for an identity certificate.

To enroll for an identity certificate using ASDM, follow these steps:

Step 1 Under **Configuration > Features > Device Administration > Certificate > Enrollment**, select the trustpoint in the **Trustpoint Name** list.

Step 2 Click **Enroll**.

After completing the action, ASDM displays the **Copy Trustpoint Configuration to Standby** dialog box, which tells you how to export the trustpoint configuration and how to check the enrollment status. This message is relevant only in a failover configuration; if you have not configured failover, you can ignore this step and click **OK**. If you have configured failover, you should follow the instructions in the dialog box to back up the certificate to the standby device.

Enrolling for a Certificate Using Manual Enrollment

Use this method when you receive an identity certificate from a CA through a means other than automatic enrollment.

Step 1 Under **Configuration > Features > Device Administration > Certificate >Trustpoint > Configuration**, click **Add**.

Step 2 On the **Add Trustpoint Configuration** dialog, type the name in the **Trustpoint Name** box.

Step 3 In the **Enrollment Settings** panel, select a key pair from the **Key Pair** list or add a new key pair by clicking **New Key Pair**.

Step 4 Optionally, type a password in the **Challenge Password** box and confirm it in the **Confirm Challenge Password** box.

Step 5 Click the **Use manual enrollment** option.

Step 6 Click **Certificate Parameters**.

a. To add subject DN values, click **Edit** in the **Certificate Parameters** dialog box.

b. In the **Edit DN** box under **DN Attribute to be Added**, select an attribute in the **Attribute** list and type a value in the **Value** box. Then click **Add**. For example, first select **Command Name (CN)** and type Pat in the **Value** box; then select **Department (OU)** and type Engineering in the **Value** box.

c. After adding all subject DN attributes, click **OK**.

d. Optionally, type values for **FQDN**, **E-mail**, and **IP Address**, and click the **Include device serial number** option.

e. Click **OK**.

Step 7 Click on **Configuration > Features > Device Administration > Certificate > Enrollment** and select the trustpoint in the **Trustpoint Name** list.

Step 8 Click **Enroll**. The **Enrollment Request** dialog box displays, which describes what to do next. After reading the instructions, click **OK**.

Either send the request by e-mail or enroll using the CA's web interface.

Step 9 After you receive the certificate from the CA, click **Configuration > Features > Device Administration > Certificate > Import Certificate** and select the name of the trustpoint in the **Trustpoint Name** list.

Step 10 Select a method for importing the certificate.

- **Import from a File**—Type the filename or browse for the file. There must be a CA certificate associated with the selected trustpoint on your system and you must have received an identity certificate in a file from the CA.
- **Enter the certificate text in base64 format—**Paste the text from the identity certificate you received from the CA into the text box. For more information, click **Help**.

Step 11 Click **Import**.

Step 12 To save the certificate enrollment configuration to flash, click **Save**.

Additional Steps for a Failover Configuration

To back up the identity certificate, CA certificate, and keys to other security appliances in your network, first export them to a file or use the export feature to display the certificate in a popup window for copying and pasting onto another security appliance through the import feature.

Exporting the Certificate to a File or PKCS12 data

To export a trustpoint configuration, follow these steps:

Step 1 Go to **Configuration > Features > Device Administration > Certificate > Trustpoint > Export**.

Step 2 Fill in the **Trustpoint Name**, **Encryption Passphrase**, and **Confirm Passphrase** fields. For information on these fields, click **Help**.

Step 3 Select a method for exporting the trustpoint configuration.

- **Export to a File**—Type the filename or browse for the file.

- **Display the trustpoint configuration in PKCS12 format—**Display the entire trustpoint configuration in a text box and then copy it for importing. For more information, click **Help**.

Step 4 Click **Export**.

Importing the Certificate onto the Standby Device

To import a trustpoint configuration, follow these steps:

Step 1 Go to **Configuration > Features > Device Administration > Certificate > Trustpoint > Import**.

Step 2 Fill in the **Trustpoint Name**, **Decryption Passphrase**, and **Confirm Passphrase** fields. For information on these fields, click **Help**. The decryption passphrase is the same as the encryption passphrase used when the trustpoint configuration was exported.

Step 3 Select a method for importing the trustpoint configuration.

- **Import from a File**—Type the filename or browse for the file.
- **Enter the trustpoint configuration in PKCS12 format—**Paste the entire trustpoint configuration from the exported source into a text box. For more information, click **Help**.

Managing Certificates

To manage certificates, go to **Configuration > Features > Device Administration > Certificate > Manage Certificates**.

You can use this panel to add a new certificate and delete a certificate. You can also display information about a certificate by clicking the **Show Detail** button.The **Certificate Details** dialog box displays three tables: **General**, **Subject** and **Issuer**.

The **General** table displays the following information:

- Type—CA, RA, or Identity.
- Serial number—Serial number of the certificate.
- Status—Available, in progress, error, fail.
- Usage—General purpose or signature.
- CRL DP—URL for of the distribution point containing the CRL for validating the certificate.
- Dates/times within which the certificate is valid— Valid from, valid to.

The **Subject** panel displays the following information:

- Name—The name of the person or entity that owns the certificate.
- Serial Number—The serial number of the security appliance.
- X.500 fields for the subject of the certificate—CN, OU, etc.
- Hostname of the certificate holder—For example, wland.com.
- Serial Number of the hostname—The serial number of the security appliance.

The **Issuer** panel displays the X.500 DN fields for the entity that granted the certificate.

- Common name (CN)
- Organizational unit or department (OU)
- Organization (O)
- Locality (L)
- State (ST)
- Country code (C)
- Email address of the issuer (EA)

CHAPTER 34

CSD

Configuration > CSD Manager

The CSD Manager window displays the following message if you choose this menu option and Cisco Secure Desktop is not installed or enabled:

```
Please install and/or enable Cisco Secure Desktop
```

Click the "Cisco Secure Desktop" link to open the Configuration >VPN > WebVPN/CSD Setup window.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

CHAPTER 35

Configuring IPS

If you are managing a Cisco ASA 5500 series adaptive security appliance equipped with an AIP SSM, you can configure the IPS features of the AIP SSM by accessing IDM from ASDM.

This section contains the following topics:

Accessing IDM from ASDM

ASDM uses IDM to configure the AIP SSM. If the AIP SSM is running IPS Version 6.0 or later, ASDM retrieves IDM from the AIP SSM and displays it as part of the ASDM interface. For earlier versions of the IPS software, IDM launches in a separate browser window.

To access IDM from ASDM, perform the following steps:

Step 1 Choose **Configuration > IPS**.

If the AIP SSM is running IPS Version 6.0 or later, ASDM retrieves IDM from the AIP SSM and displays it as part of the ASDM interface.

If the AIP SSM is running an earlier version of IPS software, ASDM displays a link to IDM.

Step 2 If a link appears in the in the ASDM pane, the AIP SSM is running a pre-6.0 version of the IPS software. Click the link to launch IDM in a new browser window. You will need to provide a username and password to access IDM.

Step 3 If a password dialog box appears, the AIP SSM is running IPS Version 6.0 or later software. Enter the AIP SSM password and click **OK**.

The IDM panes appear in the ASDM window.

If the password to access IDM is lost, you can reset the password using ASDM. See Resetting the AIP SSM Password, page 35-2, for more information.

For information about configuring IPS features, see the IDM online help. The IDM online help is available from the IDM panes displayed in ASDM. Additionally, you can see the IDM and IPS documentation on Cisco.com at the following location:

http://www.cisco.com/en/US/products/hw/vpndevc/ps4077/products_installation_and_configuration_guides_list.html

Resetting the AIP SSM Password

You can use ASDM to reset the AIP SSM password to the default if the AIP SSM is running IPS Version 6.0 or later. The default password is “cisco” (without the quotation marks). After resetting the password, you should change it to a unique value using IDM. See Accessing IDM from ASDM, page 35-1 for information about accessing IDM from ASDM.

Resetting the AIP SSM password causes the AIP SSM to reboot. IPS services are not available while the AIP SSM is rebooting.

To reset the AIP SSM password to the default, perform the following steps:

Step 1 From the ASDM menu bar, choose **Tools > IPS Password Reset**.

Note This option does not appear in the menu if an SSM is not installed. This option appears as CSC Password Reset if a CSC SSM is installed.

The IPS Password Reset confirmation dialog box appears.

Step 2 Click **OK** to reset the AIP SSM password to the default.

A dialog box displays the success or failure of the password reset. If the password was not reset, make sure you are using Version 7.2(2) or later of the platform software on the adaptive security appliance and IPS Version 6.0 or later on the AIP SSM.

Step 3 Click **Close** to close the dialog box.

CHAPTER 36

Configuring Trend Micro Content Security

ASDM lets you configure activation codes and other, basic operational parameters for the Content Security and Control (CSC) SSM as well as CSC-related features.

Managing the CSC SSM

This section contains the following topics:

About the CSC SSM

The ASA 5500 series adaptive security appliance supports the CSC SSM, which runs Content Security and Control software. The CSC SSM provides protection against viruses, spyware, spam, and other unwanted traffic. It accomplishes this by scanning the FTP, HTTP, POP3, and SMTP traffic that you configure on the adaptive security appliance to send to it.

Figure 36-1 illustrates the flow of traffic through an adaptive security appliance that has the following:

- A CSC SSM installed and configured.
- A service policy that determines what traffic is diverted to the SSM for scans.

In this example, the client could be a network user who is accessing a website, downloading files from an FTP server, or retrieving mail from a POP3 server. SMTP scans differ in that you should configure the adaptive security appliance to scan traffic sent from outside to SMTP servers protected by the adaptive security appliance.

Note The CSC SSM can scan FTP file transfers only when FTP inspection is enabled on the adaptive security appliance. By default, FTP inspection is enabled.

Figure 36-1 Flow of Scanned Traffic with CSC SSM

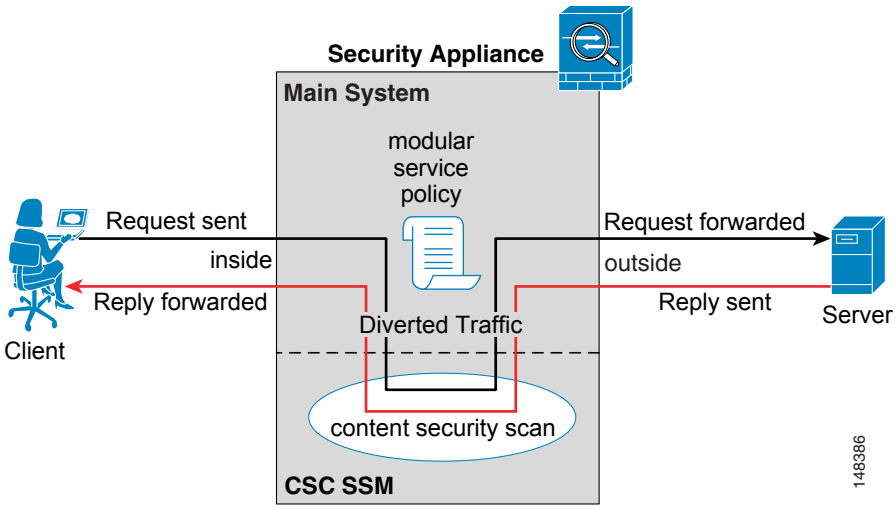

You use ASDM for system setup and monitoring of the CSC SSM. For advanced configuration of content security policies in the CSC SSM software, you access the web-based GUI for the CSC SSM by clicking links within ASDM. The CSC SSM GUI appears in a separate web browser, which may prompt you for the CSC SSM password. Use of the CSC SSM GUI is explained in the *Cisco Content Security and Control SSM Administrator Guide*.

Note ASDM and the CSC SSM maintain separate passwords. You can configure their passwords to be identical; however, changing one of these two passwords does not affect the other password.

The connection between the host running ASDM and the adaptive security appliance is made through a management port on the adaptive security appliance. The connection to the CSC SSM GUI is made through the SSM management port. Because these two connections are required to manage the CSC SSM, any host running ASDM must be able to reach the IP address of both the adaptive security appliance management port and the SSM management port.

Figure 36-2 shows an adaptive security appliance with a CSC SSM that is connected to a dedicated management network. While use of a dedicated management network is not required, we recommend it. Of particular interest in Figure 36-2 are the following:

- An HTTP proxy server is connected to the inside network and to the management network. This enables the CSC SSM to contact the Trend Micro update server.
- The management port of the adaptive security appliance is connected to the management network. To permit management of the adaptive security appliance and the CSC SSM, hosts running ASDM must be connected to the management network.
- The management network includes an SMTP server for e-mail notifications for the CSC SSM and a syslog server to which the CSC SSM can send system log messages.

Figure 36-2 CSC SSM Deployment with a Management Network

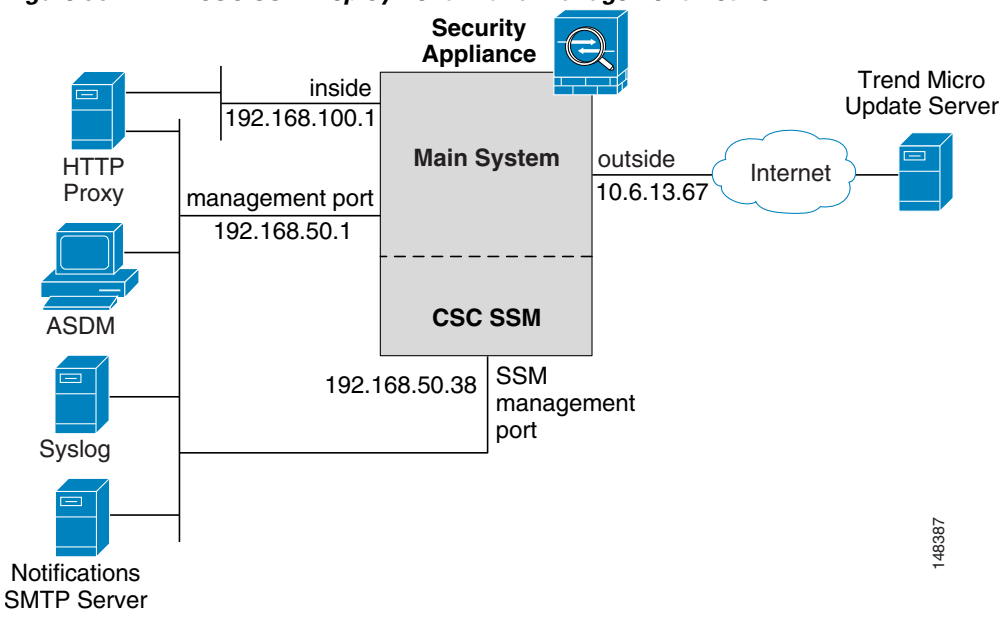

Getting Started with the CSC SSM

Before you receive the security benefits provided by a CSC SSM, you must perform several steps beyond simple hardware installation of the SSM. This procedure provides an overview of those steps.

To configure the adaptive security appliance and the CSC SSM, perform the following steps:

Step 1 If the CSC SSM did not come pre-installed in a Cisco ASA 5500 series adaptive security appliance, install it and connect a network cable to the management port of the SSM. For assistance with installation and connecting the SSM, see the *Cisco ASA 5500 Series Getting Started Guide*.

The management port of the CSC SSM must be connected to your network to allow management of and automatic updates to the CSC SSM software. Additionally, the CSC SSM uses the management port for e-mail notifications and syslogging.

Step 2 With the CSC SSM, you should have received a Product Authorization Key (PAK). Use the PAK to register the CSC SSM at the following URL:

http://www.cisco.com/go/license

After you register, you will receive activation keys by e-mail. The activation keys are required before you can complete Step 5.

Step 3 Gather the following information, for use in Step 5.

- Activation keys, received after completing Step 2.
- SSM management port IP address, netmask, and gateway IP address. The SSM management port IP address must be accessible by the hosts used to run ASDM. The IP addresses for the SSM management port and the adaptive security appliance management interface can be in different subnets.
- DNS server IP address.

- HTTP proxy server IP address (required only if your security policies require use of a proxy server for HTTP access to the Internet).
- Domain name and hostname for the SSM.
- An e-mail address and an SMTP server IP address and port number, for e-mail notifications.
- IP addresses of hosts or networks allowed to manage the CSC SSM.
- Password for the CSC SSM.

Step 4 In ASDM, verify time settings on the adaptive security appliance. Time setting accuracy is important for logging of security events and for automatic updates of CSC SSM software.

- If you manually control time settings, verify the clock settings, including time zone. Choose **Configuration > Properties > Device Administration > Clock**.
- If you are using NTP, verify the NTP configuration. Choose **Configuration > Properties > Device Administration > NTP**.

Step 5 Run the CSC Setup Wizard.

- If you have not run the CSC Setup Wizard, choose **Configuration > Trend Micro Content Security** and the Setup Wizard starts automatically.
- If you are rerunning the Setup Wizard, choose **Configuration > Trend Micro Content Security**, connect to and log into the CSC SSM, choose **CSC Setup > Wizard Setup,** and click **Launch Wizard Setup**.

For assistance with windows of the CSC Setup Wizard, click the **Help** button.

Step 6 Configure service policies to divert to the CSC SSM the traffic that you want scanned.

If you create a global service policy to divert traffic for scans, all traffic (inbound and outbound) for the supported protocols is scanned. To maximize performance of the adaptive security appliance and the CSC SSM, scan only traffic from untrusted sources.

For a discussion of best practices for diverting traffic to the CSC SSM, see Determining What Traffic to Scan.

If you want to create a global service policy that diverts traffic for scans, perform the following steps:

a. Choose **Configuration > Security Policies > Service Policy Rules** and click **Add**.

The Add Service Policy Rule Wizard appears.

b. Click the **Global - applies to all interfaces** radio button and click **Next >**.

The Traffic Classification Criteria window appears.

c. Click the **Create a new traffic class** radio button, type a name for the traffic class in the adjacent field, and check the **Any traffic** check box. Click **Next >**.

The Rules Actions window appears.

d. Click the **CSC Scan** tab and check the **Enable CSC scan for this traffic flow** check box.

e. Choose whether the adaptive security appliance should permit or deny selected traffic if the CSC SSM is unavailable by making the applicable selection in the area labeled: **If CSC card fails, then**.

f. Click **Finish**.

The new service policy appears in the Service Policy Rules pane.

g. Click **Apply**.

The adaptive security appliance begins diverting traffic to the CSC SSM, which performs the content security scans enabled by the license you purchased.

Step 7 (Optional) Review the default content security policies in the CSC SSM GUI. The default content security policies are suitable for most implementations. Modifying them is advanced configuration that you should perform only after reading the *Cisco Content Security and Control SSM Administrator Guide*.

You review the content security policies by viewing the enabled features in the CSC SSM GUI. The availability of features depends on the license level you purchased. By default, all features included in the license you purchased are enabled.

With a Base License, the features enabled by default are SMTP virus scanning, POP3 virus scanning and content filtering, webmail virus scanning, HTTP file blocking, FTP virus scanning and file blocking, logging, and automatic updates.

With a Plus License, the additional features enabled by default are SMTP anti-spam, SMTP content filtering, POP3 anti-spam, URL blocking, and URL filtering.

To access the CSC SSM GUI, in ASDM choose **Configuration > Trend Micro Content Security**, and then select one of the following: **Web**, **Mail**, **File Transfer**, or **Updates**. The blue links on these panes, beginning with the word "Configure", open the CSC SSM GUI.

Determining What Traffic to Scan

The CSC SSM can scan FTP, HTTP, POP3, and SMTP traffic. It supports these protocols only when the destination port of the packet requesting the connection is the well known port for the protocol, that is, CSC SSM can scan only the following connections:

- FTP connections opened to TCP port 21.
- HTTP connections opened to TCP port 80.
- POP3 connections opened to TCP port 110.
- SMTP connections opened to TCP port 25.

You can choose to scan traffic for all of these protocols or any combination of them. For example, if you do not allow network users to receive POP3 e-mail, you would not want to configure the adaptive security appliance to divert POP3 traffic to the CSC SSM (you would want to block it instead).

To maximize performance of the adaptive security appliance and the CSC SSM, divert to the CSC SSM only the traffic that you want the CSC SSM to scan. Needlessly diverting traffic that you do not want to scan, such as traffic between a trusted source and destination, can adversely affect network performance.

The action of scanning traffic with the CSC SSM is enabled on the CSC Scan tab of the Add Service Policy Rule Wizard—Rule Actions window. Service policies that include a CSC scan action can be applied globally or to specific interfaces; therefore, you can choose to enable CSC scans globally or for specific interfaces.

Adding the **csc** command to your global policy ensures that all unencrypted connections through the adaptive security appliance are scanned by the CSC SSM; however, this may mean that traffic from trusted sources is needlessly scanned.

If you enable CSC scans in interface-specific service policies, they are bi-directional. This means that when the adaptive security appliance opens a new connection, if a CSC scan action is active on either the inbound or the outbound interface of the connection and if the service policy identifies traffic for scanning, the adaptive security appliance diverts this traffic to the CSC SSM.

However, bi-directionality means that if you divert any of the supported traffic types that cross a given interface to the CSC SSM, it is likely performing needless scans on traffic from your trusted inside networks. For example, URLs and files requested from web servers on a DMZ network are unlikely to pose content security risks to hosts on an inside network and you probably do not want the adaptive security appliance to divert such traffic to the CSC SSM.

Therefore, we highly recommend that the service policies defining CSC scans use access lists to limit the traffic selected. Specifically, use access lists that match the following:

- HTTP connections to outside networks.
- FTP connections from clients inside the adaptive security appliance to servers outside the adaptive security appliance.
- POP3 connections from clients inside the adaptive security appliance to servers outside the adaptive security appliance.
- Incoming SMTP connections destined to inside mail servers.

In Figure 36-3, the adaptive security appliance should be configured to divert traffic to CSC SSM requests from clients on the inside network for HTTP, FTP, and POP3 connections to the outside network and incoming SMTP connections from outside hosts to the mail server on the DMZ network. HTTP requests from the inside network to the web server on the DMZ network should not be scanned.

Figure 36-3 Common Network Configuration for CSC SSM Scanning

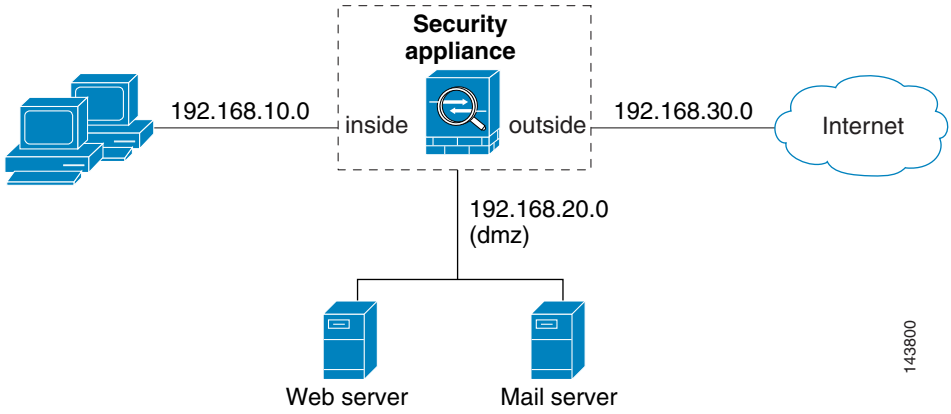

There are many ways you could configure the adaptive security appliance to identify the traffic that you want to scan. One approach is to define two service policies, one on the inside interface and the other on the outside interface, each with access lists that match traffic to be scanned.

Figure 36-4 shows service policy rules that select only the traffic that should be scanned.

Figure 36-4 Optimized Traffic Selection for CSC Scans

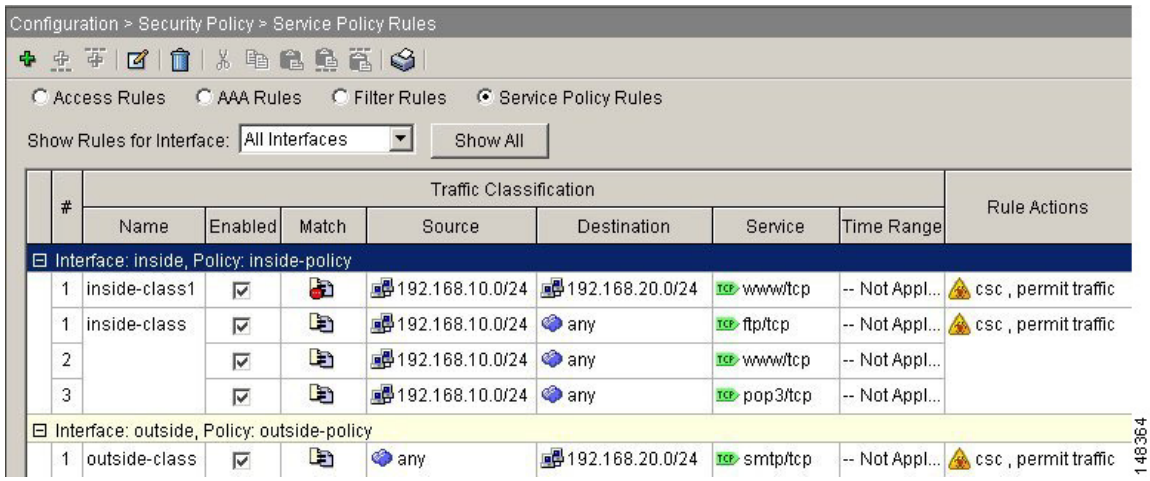

In the inside-policy, the first class, inside-class1, ensures that HTTP traffic between the inside network and the DMZ network is not scanned. The Match column indicates this by displaying the "Do not match" icon. This does not mean the adaptive security appliance blocks traffic sent from the 192.168.10.0 network to TCP port 80 on the 192.168.20.0 network. It simply exempts the traffic from being matched by the service policy applied to the inside interface and thus prevents the adaptive security appliance from sending the traffic to the CSC SSM.

The second class of the inside-policy, inside-class, matches FTP, HTTP, and POP3 traffic between the inside network and any destination. HTTP connections to the DMZ network are exempted due to inside-class1. As previously mentioned, policies applying a CSC scan action to a specific interface are effective on both ingress and egress traffic, but by specifying 192.168.10.0 as the source network, inside-class1 matches only connections initiated by the hosts on the inside network.

In the outside-policy, outside-class matches SMTP traffic from any outside source to the DMZ network. This protects the SMTP server and thus protects inside users who download e-mail from the SMTP server on the DMZ network without having to scan connections from SMTP clients to the server.

If the web server on the DMZ network receives files uploaded by HTTP from external hosts, you could add a rule to the outside policy that matches HTTP traffic from any source to the DMZ network. Because the policy is applied to the outside interface, the rule would only match connections from HTTP clients outside the adaptive security appliance.

CSC Setup

The panes under CSC Setup let you configure basic operational parameters for the CSC SSM. You must complete the Setup Wizard once before you can configure each pane separately. After you complete the Setup Wizard, you can modify each pane individually without using the Setup Wizard again.

Additionally, you cannot access the panes under Home > Content Security or Monitoring > Trend Micro Content Security until you complete the Setup Wizard. If you try to access these panes before completing the Setup Wizard, a dialog box appears and lets you access the Setup Wizard directly to complete the configuration.

For an introduction to CSC SSM, see About the CSC SSM.

- Activation/License

- IP Configuration
- Host/Notification Settings
- Management Access Host/Networks
- Password
- Restoring the Default Password
- Wizard Setup
- Summary

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	• [1]	—

1. In multiple-context mode, the panes under the CSC Setup node are available only in the admin context.

For More Information

Managing the CSC SSM

Activation/License

Configuration > Trend Micro Content Security > CSC Setup > Activation/License

The Activation/License pane lets you configure activation codes for the following two components of the CSC SSM:

- Base License
- Plus License

You can use ASDM to configure CSC licenses only once each for the two licenses. Renewed license activation codes are downloaded automatically with scheduled software updates.

Fields

- Product—*Display only.* Shows the name of the component.
- Activation Code—Contains the activation code for the corresponding Product field.
- License Status—*Display only.* Shows information about the status of the license. If the license is valid, the expiration date appears. If expiration date has passed, this field indicates that the license has expired.
- Nodes—*Display only.* Shows the maximum number of network devices supported by the Base License of your CSC SSM. The Plus License does not affect the number of network devices supported; therefore, the Nodes field does not appear in the Plus License area.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	• [1]	—

1. In multiple-context mode, the Activation/License pane is available only in the admin context.

For More Information

Managing the CSC SSM

IP Configuration

Configuration > Trend Micro Content Security > CSC Setup > IP Configuration

The IP Configuration pane lets you configure IP addresses and other relevant details for the CSC SSM, the DNS servers it should use, and a proxy server for retrieving CSC SSM software updates.

Fields

- Management Interface—Contains parameters for management access to the CSC SSM.
  - IP Address—Sets the IP address for management access to the CSC SSM.
  - Mask—Sets the netmask for the network containing the management IP address of the CSC SSM.
  - Gateway—Sets the IP address for the gateway device. This is the gateway device for the network containing the management IP address of the CSC SSM.
- DNS Servers—Contains parameters about DNS servers for the network containing the management IP address of the CSC SSM.
  - Primary DNS—Sets the IP address of the primary DNS server.
  - Secondary DNS—(Optional) Sets the IP address of the secondary DNS server.
- Proxy Server—(Optional) Contains parameters for an optional HTTP proxy server, used by the CSC SSM to contact a CSC SSM software update server. If your network configuration does not require the CSC SSM to use a proxy server, you can leave the boxes in this group empty.
  - Proxy Server—(Optional) Sets the IP address of the proxy server.
  - Proxy Port—(Optional) Sets the listening port of the proxy server.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	• [1]	—

1. In multiple-context mode, the IP Configuration pane is available only in the admin context.

For More Information

Managing the CSC SSM

Host/Notification Settings

Configuration > Trend Micro Content Security > CSC Setup > Host/Notification Settings

The Host/Notification Settings pane lets you configure details about hostname, domain name, e-mail notifications, and a domain name for e-mail messages to be excluded from detailed scanning.

Fields

- Host and Domain Names—Contains information about the hostname and domain name of the CSC SSM.
  - HostName—Sets the hostname of the CSC SSM.
  - Domain Name—Sets the domain name that contains the CSC SSM.
- Incoming E-mail Domain Name—Contains information about a trusted incoming e-mail domain name for SMTP-based e-mail.
  - Incoming Email Domain—Sets the incoming e-mail domain name. The CSC SSM scans SMTP e-mail sent to this domain. The types of threats that the CSC SSM scans for depends upon the license you purchased for the CSC SSM and the configuration of the CSC SSM software.

Note CSC SSM lets you configure a list of many incoming e-mail domains. ASDM displays only the first domain in the list. To configure additional incoming e-mail domains, access the CSC SSM interface. To do so, choose **Configuration > Trend Micro Content Security > Email** and click one of the links to access the CSC SSM. After logging in to the CSC SSM, choose **Mail (SMTP) > Configuration**, and click the **Incoming Mail** tab.

- Notification Settings—Contains information required for e-mail notification of events.
  - Administrator E-mail—Sets the e-mail address for the account to which e-mail notification should be sent.
  - E-mail Server IP Address—Sets the IP address of the SMTP server.
  - Port—Sets the port to which the SMTP server listens.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	• [1]	—

1. In multiple-context mode, the Host/Notification Settings pane is available only in the admin context.

For More Information

Managing the CSC SSM

Management Access Host/Networks

Configuration > Trend Micro Content Security > CSC Setup > Management Access Host/Networks

The Management Access Host/Networks pane lets you control the hosts and networks from which management access to the CSC SSM is permitted. You must specify at least one permitted host or network. You can specify a maximum of eight permitted hosts or networks.

Fields

- IP Address—Sets the address of a host or network you want to add to the Selected Hosts/Network list.
- Mask—Sets the netmask for the host or network you specified in the IP Address field.

  To allow all hosts and networks, enter 0.0.0.0 in the IP Address field and choose 0.0.0.0 from the Mask list.
- Selected Hosts/Networks—Displays the hosts or networks trusted for management access to the CSC SSM. ASDM requires that you configure at least one host or network. You can configure a maximum of eight hosts or networks.

  To remove a host or network from the list, choose its entry in the list and click **Delete**.
- Add >>—Adds to the Selected Hosts/Networks list the host or network you specified in the IP Address field.
- Delete—Removes the host or network selected in the Selected Hosts/Networks list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	• [1]	—

1. In multiple-context mode, the Management Access Host/Networks pane is available only in the admin context.

For More Information

Managing the CSC SSM

Password

Configuration > Trend Micro Content Security > CSC Setup > Password

The Password pane lets you change the password required for management access to the CSC SSM. The CSC SSM has a password that is maintained separately from the ASDM password. You can configure them to be identical, but changing the CSC SSM password does not affect the ASDM password.

If ASDM is connected to the CSC SSM and you change the CSC SSM password, the connection to the CSC SSM is dropped. Because of this, ASDM displays a confirmation dialog box before changing the password.

Tip Whenever the connection to the CSC SSM is dropped, you can reestablish it by using the Connection to Device icon on the status bar. To do so, click the icon and then click **Reconnect** in the Connection to Device dialog box. ASDM prompts you for the CSC SSM password, which is the new password you configured.

Note The default password is "cisco."

Passwords appear as asterisks when you type them.

Passwords must be 5 - 32 characters long.

Fields

- Old Password—Requires the current password for management access to the CSC SSM.
- New Password—Sets the new password for management access to the CSC SSM.
- Confirm New Password—Verifies the new password for management access to the CSC SSM.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	• [1]	—

1. In multiple-context mode, the Password pane is available only in the admin context.

For More Information

Managing the CSC SSM

Restoring the Default Password

Tools > CSC Password Reset

You can use ASDM to reset the CSC SSM password. You can reset this password to the default value, which is "cisco"(excluding quotation marks).

Note This option does not appear in the menu if an SSM is not installed.

To reset the CSC SSM password to the default, perform the following steps:

Step 1 From the ASDM menu bar, choose **Tools > CSC Password Reset**.

The CSC Password Reset confirmation dialog box appears.

Step 2 Click **OK** to reset the CSC SSM password to the default.

A dialog box appears, indicating the success or failure of the password reset. If the password was not reset, make sure you are using Version 7.2(2) or later on the adaptive security appliance and the most recent Version 6.1 on the CSC SSM.

Step 3 Click **Close** to close the dialog box.

Step 4 After you have reset the password, you should change it to a unique value.

Note This feature is available only in multiple-context mode in the system context.

For More Information

Password

Wizard Setup

Configuration > Trend Micro Content Security > CSC Setup > Wizard Setup

The Wizard Setup pane lets you start the Setup Wizard.

Before you can directly access any of the other panes under CSC Setup, you must complete the Setup Wizard.

After you complete the Setup Wizard, you can change any panes related to the CSC SSM without using the Setup Wizard again.

Fields

- Launch Setup Wizard—Starts the CSC SSM Setup Wizard.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	• [1]	—

1. In multiple-context mode, the Wizard pane is available only in the admin context.

For More Information

Managing the CSC SSM

Summary

Configuration > Trend Micro Content Security > CSC Setup > Wizard Setup (Summary)

The Summary pane displays the results of your actions while using the CSC Setup Wizard. The Summary pane lets you check your work before you exit the wizard. If you want to change any of the settings, you can click **< Back** to return to the panes containing those settings, make the needed changes, and click **Next >** to return to the Summary pane.

Note After you click **Finish**, you can change any panes related to the CSC SSM without using the Setup Wizard again.

Fields

- Activation Codes—*Display only.* Summarizes the settings you made on the Activation Codes Configuration pane.
  - Base—Shows the base license activation code.
  - Plus—Shows the plus license activation code, if you entered one. If not, this field is blank.
- IP Parameters—*Display only.* Summarizes the settings you made on the IP Configuration pane, including the following information:
  - IP address and netmask for the management interface of the CSC SSM.
  - IP address of the gateway device for the networking containing the CSC SSM management interface.
  - Primary DNS server IP address.
  - Secondary DNS server IP address (if configured).
  - Proxy server and port (if configured).
- Host and Domain Names—*Display only.* Summarizes the settings you made on the Host Configuration pane, including the following information:
  - Hostname of the CSC SSM.
  - Domain name for the domain containing the CSC SSM.
  - Domain name for incoming e-mail.
  - Administrator e-mail address.
  - E-mail server IP address and port number.
- Management Access List—*Display only.* Summarizes the settings you made on the Management Access Configuration pane. The drop-down list contains the hosts and networks from which the CSC SSM will allow management connections.
- Password—*Display only.* Indicates whether you changed the password on the Password Configuration pane.
- < Back—Lets you go to preceding panes of the CSC Setup Wizard.
- Next >—On the Summary pane, this button is dimmed; however, if you use Back > to access any of the preceding panes in the wizard, clicking this button returns you to the Summary pane.
- Finish—Completes the CSC Setup Wizard and saves all the settings you made while using the wizard.

- Cancel—Exits the CSC Setup Wizard without saving any of the settings you made. If you click **Cancel**, ASDM displays a dialog box to confirm your decision.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Managing the CSC SSM

Web

Configuration > Trend Micro Content Security > Web

The Web pane lets you view whether web-related features are enabled and lets you access the CSC SSM for configuring web-related features.

Note Accessing the CSC SSM requires the CSC SSM password, for which the browser displaying the CSC SSM GUI prompts you. Sessions in the CSC SSM browser window time out after ten minutes of inactivity. If you close the CSC SSM browser and click another link in ASDM, you are not prompted for the CSC SSM password because one session is already open.

Fields

- URL Blocking And Filtering—Contains information and links related to URL blocking and filtering.
  - URL Blocking—*Display only.* Shows whether the URL Blocking feature is enabled on the CSC SSM.
  - Configure URL Blocking—Opens a window for configuring URL blocking on the CSC SSM.
  - URL Filtering—*Display only.* Shows whether the URL Filtering feature is enabled on the CSC SSM.
  - Configure URL Filtering Rules—Opens a window for configuring URL filtering rules on the CSC SSM.
  - Configure URL Filtering Settings—Opens a window for configuring settings for URL filtering on the CSC SSM.
- File Blocking—Contains a field and a link about the HTTP file blocking feature on the CSC SSM.
  - File Blocking—*Display only.* Shows whether the file blocking feature is enabled on the CSC SSM.
  - Configure File Blocking—Opens a window for configuring HTTP file blocking settings on the CSC SSM.
- Scanning—Contains a field and a link about the HTTP scanning feature on the CSC SSM.

- HTTP Scanning—*Display only.* Shows whether the HTTP scanning feature is enabled on the CSC SSM.
- Configure Web Scanning—Opens a window for configuring HTTP scanning on the CSC SSM.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Managing the CSC SSM

Mail

Configuration > Trend Micro Content Security > Mail

The Mail pane lets you see if e-mail-related features are enabled and lets you access the CSC SSM for configuring these features.

For more information about configuring these areas, see the following:

- Mail > SMTP Tab
- Mail > POP3 Tab

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Mail > SMTP Tab

Configuration > Trend Micro Content Security > Mail > SMTP Tab

The SMTP tab displays fields and links specific to SMTP e-mail features on the CSC SSM.

Note Accessing the CSC SSM requires the CSC SSM password, for which the browser displaying the CSC SSM GUI prompts you. Sessions in the CSC SSM browser window time out after ten minutes of inactivity. If you close the CSC SSM browser and click another link in ASDM, you are not prompted for the CSC SSM password because one session is already open.

Fields

- Scanning—Contains fields and links about SMTP scanning.
  - Incoming Scan—*Display only.* Shows whether the incoming SMTP scanning feature is enabled on the CSC SSM.
  - Configure Incoming Scan—Opens a window for configuring incoming SMTP scan settings on the CSC SSM.
  - Outgoing Scan—*Display only.* Shows whether the outgoing SMTP scanning feature is enabled on the CSC SSM.
  - Configure Outgoing Scan—Opens a window for configuring outgoing SMTP scan settings on the CSC SSM.
- Content Filtering—Contains fields and links about SMTP content filtering.
  - Incoming Filtering—*Display only.* Shows whether content filtering for incoming SMTP e-mail is enabled on the CSC SSM.
  - Configure Incoming Filtering—Opens a window for configuring incoming SMTP content filtering settings on the CSC SSM.
  - Outgoing Filtering—*Display only.* Shows whether content filtering for outgoing SMTP e-mail is enabled on the CSC SSM.
  - Configure Outgoing Filtering—Opens a window for configuring outgoing SMTP content filtering settings on the CSC SSM.
- Anti-spam—Contains fields and links about the SMTP anti-spam feature.
  - Spam Prevention—*Display only.* Shows whether the SMTP anti-spam feature is enabled on the CSC SSM.
  - Configure Anti-spam—Opens a window for configuring SMTP anti-spam settings on the CSC SSM.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Managing the CSC SSM

Mail > POP3 Tab

Configuration > Trend Micro Content Security > Mail > POP3 Tab

The POP3 tab displays fields and links specific to POP3 e-mail features on the CSC SSM.

Note Accessing the CSC SSM requires the CSC SSM password, for which the browser displaying the CSC SSM GUI prompts you. Sessions in the CSC SSM browser window time out after ten minutes of inactivity. If you close the CSC SSM browser and click another link in ASDM, you are not prompted for the CSC SSM password because one session is already open.

Fields

- Scanning—*Display only.* Shows whether the POP3 e-mail scanning feature is enabled on the CSC SSM.
- Configure Scanning—Opens a window for configuring POP3 e-mail scanning on the CSC SSM.
- Anti-spam—*Display only.* Shows whether the POP3 anti-spam feature is enabled on the CSC SSM.
- Configure Anti-spam—Opens a window for configuring the POP3 anti-spam feature on the CSC SSM.
- Content Filtering—*Display only.* Shows whether POP3 e-mail content filtering is enabled on the CSC SSM.
- Configure Content Filtering—Opens a window for configuring POP3 e-mail content filtering on the CSC SSM.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Managing the CSC SSM

File Transfer

Configuration > Trend Micro Content Security > File Transfer

The File Transfer pane lets you view whether FTP-related features are enabled and lets you access the CSC SSM for configuring FTP-related features.

Note Accessing the CSC SSM requires the CSC SSM password, for which the browser displaying the CSC SSM GUI prompts you. Sessions in the CSC SSM browser window timeout after ten minutes of inactivity. If you close the CSC SSM browser and click another link in ASDM, you are not prompted for the CSC SSM password because one session is already open.

Fields

- File Scanning—*Display only.* Shows whether the FTP file scanning feature is enabled on the CSC SSM.

- Configure File Scanning—Opens a window for configuring FTP file scanning settings on the CSC SSM.
- File Blocking—*Display only.* Shows whether the FTP file blocking feature is enabled on the CSC SSM.
- Configure File Blocking—Opens a window for configuring FTP file blocking settings on the CSC SSM.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Managing the CSC SSM

Updates

Configuration > Trend Micro Content Security > Updates

The Updates pane lets you view whether scheduled updates are enabled and lets you access the CSC SSM for configuring scheduled updates.

Note Accessing the CSC SSM requires the CSC SSM password, for which the browser displaying the CSC SSM GUI prompts you. Sessions in the CSC SSM browser window time out after ten minutes of inactivity. If you close the CSC SSM browser and click another link in ASDM, you are not prompted for the CSC SSM password because one session is already open.

Fields

- Scheduled Updates—*Display only.* Shows whether scheduled updates are enabled on the CSC SSM.
- Scheduled Update Frequency—Displays information about when updates are scheduled to occur, such as "Hourly at 10 minutes past the hour."
- Component—Displays names of parts of the CSC SSM software that can be updated.
- Scheduled Updates—*Display only.* Shows whether scheduled updates are enabled for the corresponding components.
- Configure Updates—Opens a window for configuring scheduled update settings on the CSC SSM.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Managing the CSC SSM

Connecting to CSC/Content Security and Control Password

Home > Content Security

Configuration > Trend Micro Content Security

Monitoring > Trend Micro Content Security

With each session you start in ASDM, the first time you access features related to the CSC SSM, you must specify the management IP address and provide the password for the CSC SSM. After you successfully connect to the CSC SSM, you are not prompted again for the management IP address and password. If you start a new ASDM session, the connection to the CSC SSM is reset and you must specify the IP address and the CSC SSM password again. The connection to the CSC SSM is also reset if you change the time zone on the adaptive security appliance.

Note The CSC SSM has a password that is maintained separately from the ASDM password. You can configure them to be identical, but changing the CSC SSM password does not affect the ASDM password.

Fields

- Connecting to CSC—Lets you specify the IP address for the management port on the CSC SSM. ASDM automatically detects the IP address for the SSM in the adaptive security appliance. If this detection fails, you can specify the management IP address manually.
  - Management IP Address—Sets the management IP address for the connection to the CSC SSM to an IP address detected by ASDM. This is the default selection.
  - Other IP Address or Hostname—Sets the management IP address to the value you enter.
- CSC Password—Lets you specify the password for accessing the CSC SSM. Providing the password enables ASDM to establish a connection to the CSC SSM. It uses the connection to retrieve monitoring and status information, including information about the features enabled on the CSC SSM.

  For ten minutes after you have entered the password, clicking links that open the CSC SSM GUI does not require that you reenter the CSC SSM password in the browser that displays the CSC SSM GUI.
  - Password—Requires the CSC SSM password. If you have not completed the Setup Wizard at Configuration > Trend Micro Content Security > CSC Setup, use the default CSC password. Then complete the configuration in the Setup Wizard, which includes changing the default password.

Note The default CSC SSM password is “cisco.”

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Managing the CSC SSM

Monitoring System Log Messages

The Log viewing feature lets you view real-time system log messages that appear in the log buffer. When you open the Cisco ASDM 5.2F for FWSM window, the most recent ASDM system log messages appear at the bottom of the window. You can click the **Configure ASDM Logging Filters** link to access the Logging Filters pane. For more information about filtering system log messages, see Logging Filters, page 13-11.

You can use these messages to help troubleshoot errors or monitor system usage and performance. For a description of the Logging feature, see Chapter 13, "Configuring Logging."

About Log Viewing

This section contains the following topics:

Log Buffer

Monitoring > Logging > Log Buffer

Use this pane to view log messages saved in the buffer in a separate window.

Fields

- Logging Level—Select the level of logging messages to view, ranging from Emergency to Debugging.
- View—Opens a separate window in which log messages appear. From here you can clear the message window, and save the contents of the log. You can also search messages for specific text.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Log Buffer Viewer

Monitoring > Logging > Log Buffer > View

Use this pane to view messages that appear in the log buffer, see an explanation of the message, details about it and recommended actions to take, if necessary, to resolve it. Right-click a message in the viewer to display a menu from which you can select from the Refresh, Copy, Save, Clear, Color Settings, Create Rule, Show Rule, and Show Details options. A list of icons associated with each severity level appears at the bottom of this pane. For more information about severity levels, see Chapter 13, "Configuring Logging."

Fields

- Refresh—Refreshes the display.
- Copy—Copies a selected message.
- Save—Saves the contents of the log to your computer.
- Clear—Clears the list of messages.
- Color Settings—Enables you to specify that messages of different severity levels display in different colors.
- Create Rule—Enables you to create an access control rule that performs the opposite action of the access control rule that originally generated the message.
- Show Rule—Shows the access control rule that caused the selected message to be generated. This feature applies only to system log message IDs 106100 and 106023.
- Show Details—Shows or hides the Explanation, Recommended Action, and Details tabs. The Explanation tab provides the message syntax, an explanation for the message, and the suggested corrective action to take, if any. The Recommended Action tab describes what you should do when you receive this message. The Details tab lists the date, time, severity level, syslog ID, source IP address, destination IP address, and a description of the message.
- Find—Enter the text you want to find in the messages. Searches the messages based on the text you enter.
- Help—Provides more information.
- Filter By—Lets you enter text to filter the messages by. Press **Enter** or click **Filter** to apply the filter to the displayed messages.
- Show All—Displays all messages. Filters are removed from the display. This button is only active if a filter has been applied to the displayed log messages.
- Filter—Applies the filter to the message list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Real-Time Log Viewer

Monitoring > Logging > Real-Time Log Viewer

Use this pane to view real-time system log messages in a separate window.

Fields

- Logging Level—Select the level of logging messages to view, ranging from Emergency to Debugging.
- Buffer Limit—Maximum number of log messages to view. The default is 1000.
- View—Opens a separate window in which log messages appear. From here you can pause incoming messages, clear the message window, and save the contents of the log. You can also search messages for specific text, set color settings for different severity levels, create and show access rules, and show message details.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Real-Time Log Viewer

Monitoring > Logging > Real-Time Log Viewer >View

Use this pane to view incoming messages in real time and filter them based on text you specify. Right-click a message in the viewer to display a menu from which you can select from the Refresh, Copy, Save, Clear, Color Settings, Create Rule, Show Rule, and Show Details options. A list of icons associated with each severity level appears at the bottom of this pane. For more information about severity levels, see Chapter 13, “Configuring Logging.”

Fields

- Pause—Pauses scrolling of the Real-time Log Viewer.
- Copy—Copies a selected message.
- Save—Saves the log to your computer.
- Clear—Clears the list of messages.

- Color Settings—Enables you to specify that messages of different severity levels display in different colors.
- Create Rule—Enables you to create an access control rule that performs the opposite action of the access control rule that originally generated the message.
- Show Rule—Shows the access control rule that caused the selected message to be generated. This feature applies only to system log message IDs 106100 and 106023.
- Show Details—Shows or hides the Explanation, Recommended Action, and Details tabs. The Explanation tab provides the message syntax, an explanation for the message, and the suggested corrective action to take, if any. The Recommended Action tab describes what you should do when you receive this message. The Details tab lists the date, time, severity level, syslog ID, source IP address, destination IP address, and a description of the message.
- Find—Enter the text you want to find in the log. Searches the messages based on the text you enter.
- Help—Provides more information.
- Filter By—Lets you enter text to filter the messages by. Press **Enter** or click **Filter** to apply the filter to the displayed log messages.
- Show All—Displays all messages. Filters are removed from the display. This button is only active if a filter has been applied to the displayed log messages.
- Filter—Applies a filter to the displayed messages.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

CHAPTER 38

Monitoring Trend Micro Content Security

ASDM lets you monitor the Content Security and Control (CSC) SSM statistics as well as CSC SSM-related features.

For an introduction to CSC SSM, see About the CSC SSM.

Note If you have not completed the Setup Wizard in Configuration > Trend Micro Content Security > CSC Setup, you cannot access the panes under Monitoring > Trend Micro Content Security. Instead, a dialog box appears and lets you access the Setup Wizard directly from Monitoring > Trend Micro Content Security.

Threats

Monitoring > Trend Micro Content Security > Threats

Threats lets you view in a graph format information about various types of threats detected by the CSC SSM. You can graph a maximum of four graphs in one frame.

Fields

- Available Graphs for—Lists the components you can graph. The graphs display data in ten-second intervals.
  - Viruses detected—Displays statistics about viruses detected.
  - URL Filtered, URL Blocked—Displays statistics about URLs filtered and blocked.
  - Spam detected—Displays statistics about spam e-mail detected.
  - Spyware blocked—Displays the statistics about spyware blocked.
- Graph Window—Shows the graph window name to which you want to add a statistic type. If you have a graph window already open, a new graph window is listed by default. If you want to add a statistic type to an already open graph, choose the open graph window name. The statistics already included in the graph window are shown in the Selected Graphs list, to which you can add additional types (up to a maximum of four types per window).
- Add—Click to move the selected entries in the Available Graphs For list to the Selected Graphs list.
- Remove—Removes the selected statistic type from the Selected Graphs list.
- Show Graphs—Click to display a new or updated graph window with the selected statistics.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Managing the CSC SSM

Live Security Events

Monitoring > Trend Micro Content Security > Live Security Events

Use the Live Security Events pane to view live, real time security events in a separate window.

Fields

- Buffer Limit—The maximum number of log messages to view. The default is 1000.
- View—Opens a separate window that displays the event messages. From here you can pause incoming messages, clear the message window, and save event messages. You can also search messages for specific text.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Managing the CSC SSM

Live Security Events Viewer

Monitoring > Trend Micro Content Security > Live Security Events > Live Security Events Viewer

The Live Security Events Viewer lets you view security event messages in real time that are received from the CSC SSM. You can filter security event messages based on text you specify.

Fields

- Filter Incoming Messages

  - Show All—Displays all messages.
  - Filter by Text—Lets you filter the messages based on text you enter.
- Filter—Use to filter the messages.
- Find Messages—Searches the messages based on the text you enter.
  - Text—Enter the text to search for in the messages log.
  - Find Next—Use to find the next entry that matches the text you typed in Text.
- Columns—Displays the following, read-only columns:
  - Time—Displays the time an event occurred.
  - Source—Displays the IP address or hostname from which the threat came.
  - Threat/Filter—Displays the type of threat or, in the case of a URL filter event, the filter that triggered the event.
  - Subject/File/URL—Displays the subject of e-mails containing a threat, the names of FTP file containing a threat, or URLs blocked or filtered.
  - Receiver/Host—Displays the recipient of e-mails containing a threat or the IP address or hostname of a node threatened.
  - Sender—Displays the sender of e-mails containing a threat.
  - Content Action—Displays the action taken upon the content of a message, such as cleaning attachments or deleting attachments.
  - Msg Action—Displays the action taken upon a message, such as delivering the message unchanged, delivering the message after deleting the attachments, or delivering the message after cleaning the attachments.
- Pause—Use to pause the scrolling of the Live Security Events log.
- Save Events As—Click to save the log to your PC.
- Clear Display—Clears the list of messages.
- Close—Closes the pane and returns to the previous screen.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Managing the CSC SSM

Software Updates

Monitoring > Trend Micro Content Security > Software Updates

The Software Updates pane displays information about updates to software on the CSC SSM.

Fields

- Component—Displays names of parts of the CSC SSM software that can be updated.
- Version—Displays the current version of the corresponding component.
- Last Update—Displays the date and time that the corresponding component was updated. If the component has never been updated since the CSC SSM software was installed, “None” appears in this column.
- Last Refresh—Displays the date and time when ASDM last received information from CSC SSM regarding software updates.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Managing the CSC SSM

Resource Graphs

The security appliance lets you monitor CSC SSM status, including CPU and memory usage.

- CSC CPU
- CSC Memory

CSC CPU

Monitoring > Trend Micro Content Security > Resource Graphs > CSC CPU

The CSC CPU pane lets you view information in a graph format about CPU utilization by the CSC SSM.

Fields

- Available Graphs for—Lists the components you can graph.
  - CPU Utilization—Displays statistics for CPU use on the CSC SSM.
- Graph Window—Shows the graph window name to which you want to add a statistic type. If you have a graph window already open, a new graph window is listed by default. If you want to add a statistic type to an already open graph, choose the open graph window name. The statistics already included in the graph window are shown in the Selected Graphs list, to which you can add additional types (up to a maximum of four types per window).
- Add—Click to move the selected entries in the Available Graphs For list to the Selected Graphs list.

- Remove—Removes the selected statistic type from the Selected Graphs list.
- Show Graphs—Click to display a new or updated graph window with the selected statistics.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Managing the CSC SSM

CSC Memory

Monitoring > Trend Micro Content Security > Resource Graphs > CSC Memory

The CSC Memory pane lets you view in a graph format information about memory usage on the CSC SSM.

Fields

- Available Graphs For—Lists the components you can graph.
  - Free Memory—Displays statistics about the amount of memory not in use.
  - Used Memory—Displays statistics about the amount of memory in use.
- Graph Window—Shows the graph window name to which you want to add a statistic type. If you have a graph window already open, a new graph window is listed by default. If you want to add a statistic type to an already open graph, choose the open graph window name. The statistics already included in the graph window are shown in the Selected Graphs list, to which you can add additional types (up to a maximum of four types per window).
- Add—Click to move the selected entries in the Available Graphs For list to the Selected Graphs list.
- Remove—Removes the selected statistic type from the Selected Graphs list.
- Show Graphs—Click to display a new or updated graph window with the selected statistics.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

Managing the CSC SSM

CHAPTER 39

Monitoring Failover

Single Context Mode

Failover

You can monitor the status of the active and standby devices in a failover pair and failover related statistics. See the following screens for more information:

- Status—Displays the failover status of the device.
- Graphs—Displays graphs of various failover communication statistics.

For More Information

For more information about failover in general, see Understanding Failover.

Status

Monitoring > Properties > Failover > Status

The Status pane displays the failover state of the system. In single context mode, you can also control the failover state of the system by:

- Toggling the active/standby state of the device.
- Resetting a failed device.
- Reloading the standby unit.

Fields

Failover state of the system—*Display only.* Displays the failover state of the security appliance. The information in this field is the same output you would receive from the show failover command. The following information is included in the display:

Note Only a subset of the fields below appear when viewing the failover status within a security context. Those fields are indicated by an asterisk (*) before the field name.

- *Failover—Displays “On” when failover is enabled, “Off” when failover is not enabled.
- Cable Status—(PIX security appliance platform only) Displays the status of the serial failover cable. The following shows possible cable states:

  - Normal—The cable is connected to both units, and they both have power.
  - My side not connected—The serial cable is not connected to this unit. It is unknown if the cable is connected to the other unit.
  - Other side is not connected—The serial cable is connected to this unit, but not to the other unit.
  - Other side powered off—The other unit is turned off.
  - N/A—LAN-based failover is enabled.
- Failover unit—Displays the role of the system in the failover pair, either “Primary” or “Secondary”.
- Failover LAN Interface—Displays the logical and physical name of the LAN failover interface. If you are using the dedicated failover cable on the PIX platform, this field displays “N/A - Serial-based failover enabled”. If you have not yet configured the failover interface, this field displays “Not configured”.
- Unit Poll frequency/holdtime—Displays how often hello messages are sent on the failover link and how long to wait before testing the peer for failure if no hello messages are received.
- Interface Poll frequency—Displays the interval, in seconds, between hello messages on monitored interfaces.
- Interface Policy—Displays the number of interfaces that must fail before triggering failover.
- Monitored Interfaces—Displays the number of interfaces whose health you are monitoring for failover.
- failover replication http—Displayed if HTTP replication is enabled.
- *Last Failover—Displays the time and date the last failover occurred.
- *This Host(Context)/Other Host(Context)—For each host (or for the selected context in multiple context mode) in the failover pair, the following information is shown:
  - Primary or Secondary—Displays whether the unit is the primary or secondary unit. Also displays the following status:

    *Active—The unit is the active unit.

    *Standby—The unit is the standby unit.

    *Disabled—The unit has failover disabled or the failover link is not configured.

    *Listen—The unit is attempting to discover an active unit by listening for polling messages.

    *Learn—The unit detected an active unit, and is not synchronizing the configuration before going to standby mode.

    *Failed—The unit is failed.
  - *Active Time—The amount of time, in seconds, that the unit has been in the active state.
  - *[context_name] Interface name (n.n.n.n)—For each interface, the display shows the IP address currently being used on each unit, as well as one of the following conditions. In multiple context mode, the context name appears before each interface.

    Failed—The interface has failed.

    Link Down—The interface line protocol is down.

    Normal—The interface is working correctly.

    No Link—The interface has been administratively shut down.

    Unknown—The security appliance cannot determine the status of the interface.

    (Waiting)—The interface has not yet received any polling messages from the other unit.

Testing—The interface is being tested.

*Stateful Failover Logical Updates Statistics—The following fields relate to the Stateful Failover feature. If the Link field shows an interface name, then the Stateful Failover statistics are shown.

Note Stateful Failover is not supported on the ASA 5505 series adaptive security appliance. These statistics do not appear in ASDM running on an ASA 5505 security appliance.

- Link—Displays one of the following:
  - interface_name—The interface used for the Stateful Failover link.
  - Unconfigured—You are not using Stateful Failover.
- Stateful Obj—For each field type, the following statistics are displayed:

  xmit—Number of transmitted packets to the other unit

  xerr—Number of errors that occurred while transmitting packets to the other unit

  rcv—Number of received packets

  rerr—Number of errors that occurred while receiving packets from the other unit

  The following are the stateful object field types:
  - General—Sum of all stateful objects.
  - sys cmd—Logical update system commands; for example, LOGIN and Stay Alive.
  - up time—Up time, which the active unit passes to the standby unit.
  - RPC services—Remote Procedure Call connection information.
  - TCP conn—TCP connection information.
  - UDP conn—Dynamic UDP connection information.
  - ARP tbl—Dynamic ARP table information.
  - L2BRIDGE tbl—Layer 2 bridge table information (transparent firewall mode only).
  - Xlate_Timeout—Indicates connection translation timeout information.
  - VPN IKE upd—IKE connection information.
  - VPN IPSEC upd—IPSec connection information.
  - VPN CTCP upd—cTCP tunnel connection information.
  - VPN SDI upd—SDI AAA connection information.
  - VPN DHCP upd—Tunneled DHCP connection information.
- *Logical Update Queue Information—Displays the following statistics:
  - Recv Q—The status of the receive queue.
  - Xmit Q—The status of the transmit queue.

  The following information is displayed for each queue:
  - Cur—The current number of packets in the queue.
  - Max—The maximum number of packets.
  - Total—The total number of packets.

*Lan-based Failover is active—This field appears only when LAN-based failover is enabled.

- interface name (n.n.n.n) and peer (n.n.n.n)—The name and IP address of the failover link currently being used on each unit.

The following actions are available on the Status pane:

- Make Active—(Only available in Single mode) Click this button to make the security appliance the active unit in an active/standby configuration.
- Make Standby—(Only available in Single mode) Click this button to make the security appliance the standby unit in an active/standby pair.
- Reset Failover—(Only available in Single mode) Click this button to reset a system from the failed state to the standby state. You cannot reset a system to the active state. Clicking this button on the active unit resets the standby unit.
- Reload Standby—(Only available in Single mode) Click this button to force the standby unit to reload.
- Refresh—Click this button to refresh the status information in the Failover state of the system field.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

For more information about failover in general, see Understanding Failover.

Graphs

Monitoring > Properties > Failover > Graphs

The Graphs pane lets you view failover statistics in graph and table form. In multiple context mode, the Graphs pane is only available in the admin context.

The information in the graphs relate to Stateful Failover only.

Fields

- Available Graphs for—Lists the types of statistical information available for monitoring. You can choose up to four statistic types to display in one graph window. Double-clicking a statistic type in this field moves it to the Selected Graphs field. Single-clicking a statistic type in this field selects the entry. You can select multiple entries.

  The following types of statistics are available in graph or table format in the graph window. They show the number of packets sent to and received from the other unit in the failover pair.

  - RPC services information—Displays the security appliance RPC service information.
  - TCP Connection Information—Displays the security appliance TCP connection information.
  - UDP Connection Information—Displays the security appliance UDP connection information.
  - ARP Table Information—Displays the security appliance ARP table information.

  - L2Bridge Table Information—(Transparent Firewall Mode Only) Displays the layer 2 bridge table packet counts.
  - Xmit Queue—(Single Mode Only) Displays the current, maximum, and total number of packets transmitted.
  - Receive Queue—(Single Mode Only) Displays the current, maximum, and total number of packets received.

- Graph Window—Shows the graph window name to which you want to add a statistic type. If you have a graph window already open, a new graph window is listed by default. If you want to add a statistic type to an already open graph, select the open graph window name. The statistics already included in the graph window are shown in the Selected Graphs field, to which you can add additional types (up to a maximum of four types per window).
- Add—Click this button to move the selected entries in the Available Graphs for field to the Selected Graphs field.
- Remove—Removes the selected statistic type from the Selected Graphs field.
- Selected Graphs—Shows the statistic types you want to show in the selected graph window. You can include up to four types. Double-clicking a statistic type in this field removes the selected statistic type from the field. Single-clicking a statistic type in this field selects the statistic type. You can select multiple statistic types.
- Show Graphs—Click this button to display a new or updated graph window with the selected statistics.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

For More Information

For more information about failover in general, see Understanding Failover.

Multiple Context Mode

You can monitor the failover status of the system and of the individual failover groups in the system context. See the following topics for monitoring failover status from the system context:

- System
- Failover Group 1 and Failover Group 2

For More Information

For more information about failover in general, see Understanding Failover.

System

System > Monitoring > Failover > System

The System pane displays the failover state of the system.You can also control the failover state of the system by:

- Toggling the active/standby state of the device.
- Resetting a failed device.
- Reloading the standby unit.

Fields

Failover state of the system—*Display only.* Displays the failover state of the security appliance. The information shown is the same output you would receive from the **show failover** command. The following information is included in the display:

- Failover—Displays “On” when failover is enabled, “Off” when failover is not enabled.
- Cable Status—(PIX security appliance platform only) Displays the status of the serial failover cable. The following shows possible cable states:
  - Normal—The cable is connected to both units, and they both have power.
  - My side not connected—The serial cable is not connected to this unit. It is unknown if the cable is connected to the other unit.
  - Other side is not connected—The serial cable is connected to this unit, but not to the other unit.
  - Other side powered off—The other unit is turned off.
  - N/A—LAN-based failover is enabled.
- Failover unit—Displays the role of the system in the failover pair, either “Primary” or “Secondary”.
- Failover LAN Interface—Displays the logical and physical name of the LAN failover interface. If you are using the dedicated failover cable on the PIX platform, this field displays “N/A - Serial-based failover enabled”. If you have not yet configured the failover interface, this field displays “Not configured”.
- Unit Poll frequency/holdtime—Displays how often hello messages are sent on the failover link and how long to wait before testing the peer for failure if no hello messages are received.
- Interface Poll frequency—Displays the interval, in seconds, between hello messages on monitored interfaces.
- Interface Policy—Displays the number of interfaces that must fail before triggering failover.
- Monitored Interfaces—Displays the number of interfaces whose health you are monitoring for failover.
- failover replication http—Specifies that HTTP replication is enabled.
- Group x Last Failover—Displays the time and date the last failover occurred for each failover group.
- This Host/Other Host —For each host in the failover pair, the following information is shown:
  - Primary or Secondary—Displays whether the unit is the primary or secondary unit.
  - Group x—For each failover group, the following information is shown:

    State—Active or Standby Ready.

Active Time—The amount of time, in seconds, that the failover group has been in the active state.

- context_name Interface name (n.n.n.n)—For each interface, the display shows the IP address currently being used on each unit, as well as one of the following conditions.

  Failed—The interface has failed.

  Link Down—The interface line protocol is down.

  Normal—The interface is working correctly.

  No Link—The interface has been administratively shut down.

  Unknown—The security appliance cannot determine the status of the interface.

  (Waiting)—The interface has not yet received any polling messages from the other unit.

  Testing—The interface is being tested.

Stateful Failover Logical Updates Statistics—The following fields relate to the Stateful Failover feature. If the Link field shows an interface name, then the Stateful Failover statistics are shown.

Note Stateful Failover is not supported on the ASA 5505 series adaptive security appliance. These statistics do not appear in ASDM running on an ASA 5505 security appliance.

- Link—Displays one of the following:
  - *interface_name*—The interface used for the Stateful Failover link.
  - Unconfigured—You are not using Stateful Failover.
- Stateful Obj—For each field type, the following statistics are displayed:

  xmit—Number of transmitted packets to the other unit

  xerr—Number of errors that occurred while transmitting packets to the other unit

  rcv—Number of received packets

  rerr—Number of errors that occurred while receiving packets from the other unit

  The following are the stateful object field types:
  - General—Sum of all stateful objects.
  - sys cmd—Logical update system commands; for example, LOGIN and Stay Alive.
  - up time—Up time, which the active unit passes to the standby unit.
  - RPC services—Remote Procedure Call connection information.
  - TCP conn—TCP connection information.
  - UDP conn—Dynamic UDP connection information.
  - ARP tbl—Dynamic ARP table information.
  - L2BRIDGE tbl—Layer 2 bridge table information (transparent firewall mode only).
  - Xlate_Timeout—Indicates connection translation timeout information.
  - VPN IKE upd—IKE connection information.
  - VPN IPSEC upd—IPSec connection information.
  - VPN CTCP upd—cTCP tunnel connection information.
  - VPN SDI upd—SDI AAA connection information.

  - VPN DHCP upd—Tunneled DHCP connection information.
- Logical Update Queue Information—Displays the following statistics:
  - Recv Q—The status of the receive queue.
  - Xmit Q—The status of the transmit queue.

  The following information is displayed for each queue:
  - Cur—The current number of packets in the queue.
  - Max—The maximum number of packets.
  - Total—The total number of packets.

Lan-based Failover is active—This field appears only when LAN-based failover is enabled.

- interface *name* (*n.n.n.n*) and peer (*n.n.n.n*)—The name and IP address of the failover link currently being used on each unit.

The following actions are available on the System pane:

- Make Active—Click this button to make the security appliance the active unit in an active/standby configuration. In an active/active configuration, clicking this button causes both failover groups to become active on the security appliance.
- Make Standby—Click this button to make the security appliance the standby unit in an active/standby pair. In an active/active configuration, clicking this button causes both failover groups to go to the standby state on the security appliance.
- Reset Failover—Click this button to reset a system from the failed state to the standby state. You cannot reset a system to the active state. Clicking this button on the active unit resets the standby unit.
- Reload Standby—Click this button to force the standby unit to reload.
- Refresh—Click this button to refresh the status information in the Failover state of the system field.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

For More Information

For more information about failover in general, see Understanding Failover.

Failover Group 1 and Failover Group 2

System > Monitoring > Failover > Failover Group 1 and Failover Group 2

The Failover Group 1 and Failover Group 2 panes display the failover state of the selected group.You can also control the failover state of the group by toggling the active/standby state of the group or by resetting a failed group.

Fields

Failover state of Group[*x*]—*Display only*. Displays the failover state of the selected failover group. The information shown is the same as the output you would receive from the **show failover group** command and contains the following information:

- Last Failover—The time and date of the last failover.
- This Host/Other Host—For each host in the failover pair, the following information is shown:
  - Primary or Secondary—Displays whether the unit is the primary or secondary unit. The following information is also shown for the failover group:

    Active—The failover group is active on the specified unit.

    Standby—The failover group is in the standby state on the specified unit.

    Disabled—The unit has failover disabled or the failover link is not configured.

    Listen—The unit is attempting to discover an active unit by listening for polling messages.

    Learn—The unit detected an active unit, and is not synchronizing the configuration before going to standby mode.

    Failed—The failover group is in the failed state on the specified unit.
  - Active Time—The amount of time, in seconds, that the failover group has been in the active state on the specified unit.
  - *context_name* Interface *name* (n.n.n.n)—For each interface in the selected failover group, the display shows the context to which it belongs and the IP address currently being used on each unit, as well as one of the following conditions.

    Failed—The interface has failed.

    Link Down—The interface line protocol is down.

    Normal—The interface is working correctly.

    No Link—The interface has been administratively shut down.

    Unknown—The security appliance cannot determine the status of the interface.

    (Waiting)—The interface has not yet received any polling messages from the other unit.

    Testing—The interface is being tested.
- Stateful Failover Logical Updates Statistics—The following fields relate to the Stateful Failover feature. If the Link field shows an interface name, then the Stateful Failover statistics are shown.

  Link—Displays one of the following:
  - *interface_name*—The interface used for the Stateful Failover link.
  - Unconfigured—You are not using Stateful Failover.

  Stateful Obj—For each field type, the following statistics are displayed:
  - xmit—Number of transmitted packets to the other unit
  - xerr—Number of errors that occurred while transmitting packets to the other unit
  - rcv—Number of received packets
  - rerr—Number of errors that occurred while receiving packets from the other unit

  The following are the stateful object field types:
  - General—Sum of all stateful objects.
  - sys cmd—Logical update system commands; for example, LOGIN and Stay Alive.

  - up time—Up time, which the active unit passes to the standby unit.
  - RPC services—Remote Procedure Call connection information.
  - TCP conn—TCP connection information.
  - UDP conn—Dynamic UDP connection information.
  - ARP tbl—Dynamic ARP table information.
  - L2BRIDGE tbl—Layer 2 bridge table information (transparent firewall mode only).
  - Xlate_Timeout—Indicates connection translation timeout information.
  - IKE upd—IKE connection information.
  - VPN IPSEC upd—IPSec connection information.
  - VPN CTCP upd—cTCP tunnel connection information.
  - VPN SDI upd—SDI AAA connection information.
  - VPN DHCP upd—Tunneled DHCP connection information.
- Logical Update Queue Information—Displays the following statistics:
  - Recv Q—The status of the receive queue.
  - Xmit Q—The status of the transmit queue.

  The following information is displayed for each queue:
  - Cur—The current number of packets in the queue.
  - Max—The maximum number of packets.
  - Total—The total number of packets.

You can performthe following actions from this pane:

- Make Active—Click this button to make the failover group active unit on the security appliance.
- Make Standby—Click this button to force the failover group into the standby state on the security appliance.
- Reset Failover—Click this button to reset a system from the failed state to the standby state. You cannot reset a system to the active state. Clicking this button on the active unit resets the standby unit.
- Refresh—Click this button to refresh the status information in the Failover state of the system field.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	—	—	•

For More Information

For more information about failover in general, see Understanding Failover.

Monitoring Interfaces

ASDM lets you monitor interface statistics as well as interface-related features.

ARP Table

Monitoring > Interfaces > ARP Table

The ARP Table pane displays the ARP table, including static and dynamic entries. The ARP table includes entries that map a MAC address to an IP address for a given interface. See Configuration > Properties > ARP Static Table for more information about the ARP table.

Fields

- Interface—Lists the interface name associated with the mapping.
- IP Address—Shows the IP address.
- MAC Address—Shows the MAC address.
- Proxy ARP—Displays Yes if proxy ARP is enabled on the interface. Displays No if proxy ARP is not enabled on the interface.
- Clear—Clears the dynamic ARP table entries. Static entries are not cleared.
- Refresh—Refreshes the table with current information from the security appliance and updates Last Updated date and time.
- Last Updated—*Display only.* Shows the date and time the display was updated.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

DHCP

The security appliance lets you monitor DHCP status, including the addresses assigned to clients, the lease information for a security appliance interface, and DHCP statistics.

DHCP Server Table

Monitoring > Interfaces > DHCP > DHCP Server Table

The DHCP Server Table lists the IP addresses assigned to DHCP clients.

Fields

- IP Address—Shows the IP address assigned to the client.
- Client-ID—Shows the client MAC address or ID.
- Lease Expiration—Shows the date that the DHCP lease expires. The lease indicates how long the client can use the assigned IP address. Remaining time is also specified in the number of seconds and is based on the timestamp in the Last Updated display-only field.
- Number of Active Leases—Shows the total number of DHCP leases.
- Refresh—Refreshes the information from the security appliance.
- Last Updated—Shows when the data in the table was last updated.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

DHCP Client Lease Information

Monitoring > Interfaces > DHCP > DHCP Server Table

If you obtain the security appliance interface IP address from a DHCP server, the DHCP Client Lease Information panel shows information about the DHCP lease.

Fields

- Select an interface—Lists the security appliance interfaces. Choose the interface for which you want to view the DHCP lease. If an interface has multiple DHCP leases, then choose the interface and IP address pair you want to view.
- Attribute and Value—Lists the attributes and values of the interface DHCP lease.
  - Temp IP addr—*Display only.* The IP address assigned to the interface.
  - Temp sub net mask—*Display only.* The subnet mask assigned to the interface.
  - DHCP lease server—*Display only.* The DHCP server address.

- state—*Display only.* The state of the DHCP lease, as follows:

  Initial—The initialization state, where the security appliance begins the process of acquiring a lease. This state is also shown when a lease ends or when a lease negotiation fails.

  Selecting—The security appliance is waiting to receive DHCPOFFER messages from one or more DHCP servers, so it can choose one.

  Requesting—The security appliance is waiting to hear back from the server to which it sent its request.

  Purging—The security appliance is removing the lease because of an error.

  Bound—The security appliance has a valid lease and is operating normally.

  Renewing—The security appliance is trying to renew the lease. It regularly sends DHCPREQUEST messages to the current DHCP server, and waits for a reply.

  Rebinding—The security appliance failed to renew the lease with the original server, and now sends DHCPREQUEST messages until it gets a reply from any server or the lease ends.

  Holddown—The security appliance started the process to remove the lease.

  Releasing—The security appliance sends release messages to the server indicating that the IP address is no longer needed.
- Lease—*Display only.* The length of time, specified by the DHCP server, that the interface can use this IP address.
- Renewal—*Display only.* The length of time until the interface automatically attempts to renew this lease.
- Rebind—*Display only.* The length of time until the security appliance attempts to rebind to a DHCP server. Rebinding occurs if the security appliance cannot communicate with the original DHCP server, and 87.5 percent of the lease time has expired. The security appliance then attempts to contact any available DHCP server by broadcasting DHCP requests.
- Next timer fires after—*Display only.* The number of seconds until the internal timer triggers.
- Retry count—*Display only.* If the security appliance is attempting to establish a lease, this field shows the number of times the security appliance tried sending a DHCP message. For example, if the security appliance is in the Selecting state, this value shows the number of times the security appliance sent discover messages. If the security appliance is in the Requesting state, this value shows the number of times the security appliance sent request messages.
- Client-ID—*Display only.* The client ID used in all communication with the server.
- Proxy—*Display only.* Specifies if this interface is a proxy DHCP client for VPN clients, True or False.
- Hostname—*Display only.* The client hostname.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

DHCP Statistics

Monitoring > Interfaces > DHCP > DHCP Statistics

The DHCP Statistics pane shows statistics for the DHCP server feature.

Fields

- Message Type—Lists the DHCP message types sent or received:
  - BOOTREQUEST
  - DHCPDISCOVER
  - DHCPREQUEST
  - DHCPDECLINE
  - DHCPRELEASE
  - DHCPINFORM
  - BOOTREPLY
  - DHCPOFFER
  - DHCPACK
  - DHCPNAK
- Count—Shows the number of times a specific message was processed.
- Direction—Shows if the message type is Sent or Received.
- Total Messages Received—Shows the total number of messages received by the security appliance.
- Total Messages Sent—Shows the total number of messages sent by the security appliance.
- Counter—Shows general statistical DHCP data, including the following:
  - DHCP UDP Unreachable Errors
  - DHCP Other UDP Errors
  - Address Pools
  - Automatic Bindings
  - Expired Bindings
  - Malformed Messages
- Value—Shows the number of each counter item.
- Refresh—Updates the DHCP table listings.
- Last Updated—Shows when the data in the tables was last updated.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

MAC Address Table

Monitoring > Interfaces > MAC Address Table

The MAC Address Table pane shows the static and dynamic MAC address entries. See Configuration > Properties > Bridging > MAC Address Table for more information about the MAC address table and adding static entries.

Fields

- Interface—Shows the interface name associated with the entry.
- MAC Address—Shows the MAC address.
- Type—Shows if the entry is static or dynamic.
- Age—Shows the age of the entry, in minutes. To set the timeout, see MAC Address Table.
- Refresh—Refreshes the table with current information from the security appliance.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
—	•	•	•	—

Dynamic ACLs

Monitoring > Interfaces > Dynamic ACLs

The Dynamic ACLs pane shows a table of the Dynamic ACLs, which are functionally identical to the user-configured ACLs except that they are created, activated and deleted automatically by the security appliance. These ACLs do not show up in the configuration and are only visible in this table. They are identified by the "(dynamic)" keyword in the ACL header.

When you choose an ACL in this table, the contents of the ACL is shown in the bottom text field.

Fields

- ACL—Shows the name of the dynamic ACL.
- Element Count—Shows the number of elements in the ACL
- Hit Count—Shows the total hit count for all of the elements in the ACL.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Interface Graphs

Monitoring > Interfaces > Interface Graphs

The Interface Graphs pane lets you view interface statistics in graph or table form. If an interface is shared among contexts, the security appliance shows only statistics for the current context. The number of statistics shown for a subinterface is a subset of the number of statistics shown for a physical interface.

Fields

- Available Graphs for—Lists the types of statistics available for monitoring. 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.
  - Byte Counts—Shows the number of bytes input and output on the interface.
  - Packet Counts—Shows the number of packets input and output on the interface.
  - Packet Rates—Shows the rate of packets input and output on the interface.
  - Bit Rates—Shows the bit rate for the input and output of the interface.
  - Drop Packet Count—Shows the number of packets dropped on the interface.

  These additional statistics display for physical interfaces:
  - Buffer Resources—Shows the following statistics:

    Overruns—The number of times that the security appliance was incapable of handing received data to a hardware buffer because the input rate exceeded the security appliance capability to handle the data.

    Underruns—The number of times that the transmitter ran faster than the security appliance could handle.

    No Buffer—The number of received packets discarded because there was no buffer space in the main system. Compare this with the ignored count. Broadcast storms on Ethernet networks are often responsible for no input buffer events.
  - Packet Errors—Shows the following statistics:

    CRC—The number of Cyclical Redundancy Check errors. When a station sends a frame, it appends a CRC to the end of the frame. This CRC is generated from an algorithm based on the data in the frame. If the frame is altered between the source and destination, the security appliance notes that the CRC does not match. A high number of CRCs is usually the result of collisions or a station transmitting bad data.

    Frame—The number of frame errors. Bad frames include packets with an incorrect length or bad frame checksums. This error is usually the result of collisions or a malfunctioning Ethernet device.

Input Errors—The number of total input errors, including the other types listed here. Other input-related errors can also cause the input error count to increase, and some datagrams might have more than one error; therefore, this sum might exceed the number of errors listed for the other types.

Runts—The number of packets that are discarded because they are smaller than the minimum packet size, which is 64 bytes. Runts are usually caused by collisions. They might also be caused by poor wiring and electrical interference.

Giants—The number of packets that are discarded because they exceed the maximum packet size. For example, any Ethernet packet that is greater than 1518 bytes is considered a giant.

Deferred—For FastEthernet interfaces only. The number of frames that were deferred before transmission due to activity on the link.

- Miscellaneous—Shows statistics for received broadcasts.
- Collision Counts—For FastEthernet interfaces only. Shows the following statistics:

  Output Errors—The number of frames not transmitted because the configured maximum number of collisions was exceeded. This counter should only increment during heavy network traffic.

  Collisions—The number of messages retransmitted due to an Ethernet collision (single and multiple collisions). This usually occurs on an overextended LAN (Ethernet or transceiver cable too long, more than two repeaters between stations, or too many cascaded multiport transceivers). A packet that collides is counted only once by the output packets.

  Late Collisions—The number of frames that were not transmitted because a collision occurred outside the normal collision window. A late collision is a collision that is detected late in the transmission of the packet. Normally, these should never happen. When two Ethernet hosts try to talk at once, they should collide early in the packet and both back off, or the second host should see that the first one is talking and wait. If you get a late collision, a device is jumping in and trying to send the packet on the Ethernet while the security appliance is partly finished sending the packet. The security appliance does not resend the packet, because it may have freed the buffers that held the first part of the packet. This is not a real problem because networking protocols are designed to cope with collisions by resending packets. However, late collisions indicate a problem exists in your network. Common problems are large repeated networks and Ethernet networks running beyond the specification.

- Input Queue—Shows the number of packets in the input queue, the current and the maximum, including the following statistics:

  Hardware Input Queue—The number of packets in the hardware queue.

  Software Input Queue—The number of packets in the software queue.

- Output Queue—Shows the number of packets in the output queue, the current and the maximum, including the following statistics:

  Hardware Output Queue—The number of packets in the hardware queue.

  Software Output Queue—The number of packets in the software queue.

- Drop Packet Queue—Shows the number of packets dropped.

- Add—Adds the selected statistic type to the selected graph window.
- Remove—Removes the selected statistic type from the selected graph window. This button name changes to Delete if the item you are removing was added from another panel, and is not being returned to the Available Graphs pane.

- Show Graphs—Shows the graph window name to which you want to add a statistic type. If you have a graph window already open, a new graph window is listed by default. If you want to add a statistic type to an already open graph, choose the open graph window name. The statistics already included on the graph are shown in the Selected Graphs pane, to which you can add additional types. Graph windows are named for ASDM followed by the interface IP address and the name "Graph". Subsequent graphs are named "Graph (2)" and so on.
- Selected Graphs—Shows the statistic types you want to show in the selected graph window. You an include up to four types.
  - Show Graphs—Shows the graph window or updates the graph with additional statistic types if added.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Graph/Table

Monitoring > Interfaces > Interface Graphs > Graph/Table

The Graph window shows a graph for the selected statistics. The Graph window can show up to four graphs and tables at a time. By default, the graph or table displays the real-time statistics. If you enable History Metrics, you can view statistics for past time periods.

Fields

- View—Sets the time period for the graph or table. To view any time period other than real-time, enable History Metrics. The data is updated according to the specification of the following options:
  - 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 min
  - Last 5 days, data every 2 hours
- Export—Exports the graph in comma-separated value format. If there is more than one graph or table on the Graph window, the Export Graph Data dialog box appears. Choose one or more of the graphs and tables listed by checking the box next to the name.
- Print—Prints the graph or table. If there is more than one graph or table on the Graph window, the Print Graph dialog box appears. Choose the graph or table you want to print from the Graph/Table Name list.
- Bookmark—Opens a browser window with a single link for all graphs and tables on the Graphs window, as well as individual links for each graph or table. You can then copy these URLs as bookmarks in your browser. ASDM does not have to be running when you open the URL for a graph; the browser launches ASDM and then displays the graph.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

PPPoE Client

Monitoring > Interfaces > PPPoE Client

The PPPoE Client Lease Information pane displays information about current PPPoE connections.

Fields

Select a PPPoE interface—Select an interface that you want to view PPPoE client lease information.

Refresh—loads the latest PPPoE connection information from the security appliance for display.

interface connection

Monitoring > Interfaces > *interface* connection

The *interface* connection node in the Monitoring > Interfaces tree only appears if static route tracking is configured. If you have several routes tracked, there will be a node for each interface that contains a tracked route.

See the following for more information about the route tracking information available:

- Track Status for, page 40-9
- Monitoring Statistics for, page 40-10

Track Status for

Monitoring > Interfaces > *interface* connection > Track Status for

The Track Status for pane displays information about the the tracked object.

Fields

- Tracked Route—*Display only.* Displays the route associated with the tracking process.
- Route Statistics—*Display only.* Displays the reachability of the object, when the last change in reachability occurred, the operation return code, and the process that is performing the tracking.

Modes

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Monitoring Statistics for

Monitoring > Interfaces > *interface* connection > Monitoring Statistics for

The Monitoring Statics for pane displays statistics for the SLA monitoring process.

Fields

- SLA Monitor ID—*Display only.* Displays the ID of the SLA monitoring process.
- SLA statistics—*Display only.* Displays SLA monitoring statistics, such as the last time the process was modified, the number of operations attempted, the number of operations skipped, and so on.

Modes

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

CHAPTER 41

Monitoring Routing

You can monitor the following routing information on the security appliance:

- OSPF LSAs
- OSPF Neighbors
- Routes

OSPF LSAs

You can view the LSAs stored in the security appliance OSPF database. There are 4 types of LSAs stored in the database, each with its own particular format. The following briefly describes the LSA types:

- Router LSAs (Type 1 LSAs) describe the routers attached to a network.
- Network LSAs (Type 2 LSAs) describe the networks attached to an OSPF router.
- Summary LSAs (Type 3 and Type 4 LSAs) condense routing information at area borders.
- External LSAs (Type 5 and Type 7 LSAs) describe routes to external networks.

To learn more about the information displayed for each LSAs type, see the following:

- Type 1
- Type 2
- Type 3
- Type 4
- Type 5
- Type 7

Type 1

Monitoring > Routing > Routing > OSPF LSAs > Type 1

Type 1 LSAs are router link advertisements that are passed within an area by all OSPF routers. They describe the router links to the network. Type 1 LSAs are only flooded within a particular area.

The Type 1 pane displays all Type 1 LSAs received by the security appliance. Each row in the table represents a single LSA.

Fields

- Process—*Display only.* Displays the OSPF process for the LSA.
- Area—*Display only.* Displays the OSPF area for the LSA.
- Router ID—*Display only.* Displays the OSPF router ID of the router originating the LSA.
- Advertiser—*Display only.* Displays the ID of the router originating the LSA. For router LSAs, this is identical to the Router ID.
- Age—*Display only.* Displays the age of the link state.
- Sequence #—*Display only.* Displays the link state sequence number. The link state sequence number is used to detect old or duplicate LSAs.
- Checksum—*Display only.* Displays the checksum of the contents of the LSA.
- Link Count—*Display only.* Displays the number of interfaces detected for the router.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Type 2

Monitoring > Routing > Routing > OSPF LSAs > Type 2

Type 2 LSAs are network link advertisements that are flooded within an area by the Designated Router. They describe the routers attached to specific networks.

The Type 2 pane displays the IP address of the Designated Router that advertises the routes.

Fields

- Process—*Display only.* Displays the OSPF process for the LSA.
- Area—*Display only.* Displays the OSPF area for the LSA.
- Designated Router—*Display only.* Displays the IP address of the Designated Router interface that sent the LSA.
- Advertiser—*Display only.* Displays the OSPF router ID of the Designated Router that sent the LSA.
- Age—*Display only.* Displays the age of the link state.
- Sequence #—*Display only.* Displays the link state sequence number. The link state sequence number is used to detect old or duplicate LSAs.
- Checksum—*Display only.* Displays the checksum of the contents of the LSA.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Type 3

Monitoring > Routing > Routing > OSPF LSAs > Type 3

Type 3 LSA are summary link advertisements that are passed between areas. They describe the networks within an area.

Fields

- Process—*Display only.* Displays the OSPF process for the LSA.
- Area—*Display only.* Displays the OSPF area for the LSA.
- Destination—*Display only.* Displays the address of the destination network being advertised.
- Advertiser—*Display only.* Displays the ID of the ABR that sent the LSA.
- Age—*Display only.* Displays the age of the link state.
- Sequence #—*Display only.* Displays the link state sequence number. The link state sequence number is used to detect old or duplicate LSAs.
- Checksum—*Display only.* Displays the checksum of the contents of the LSA.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Type 4

Monitoring > Routing > Routing > OSPF LSAs > Type 3

Type 4 LSAs are summary link advertisements that are passed between areas. They describe the path to the ASBR. Type 4 LSAs do not get flooded into stub areas.

Fields

- Process—*Display only.* Displays the OSPF process for the LSA.
- Area—*Display only.* Displays the OSPF area for the LSA.
- Router ID—*Display only.* Displays the router ID of the ASBR.
- Advertiser—*Display only.* Displays the ID of the ABR that sent the LSA.

- Age—*Display only.* Displays the age of the link state.
- Sequence #—*Display only.* Displays the link state sequence number. The link state sequence number is used to detect old or duplicate LSAs.
- Checksum—*Display only.* Displays the checksum of the contents of the LSA.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Type 5

Monitoring > Routing > Routing > OSP LSAs > Type 3

Type 5 LSAs are passed between and flooded into areas by ABSRs. They describe routes external to the AS. Stub areas and NSSAs do not receive these LSAs.

Fields

- Process—*Display only.* Displays the OSPF process for the LSA.
- Network—*Display only.* Displays the address of the AS external network.
- Advertiser—*Display only.* Displays the router ID of the ASBR.
- Age—*Display only.* Displays the age of the link state.
- Sequence #—*Display only.* Displays the link state sequence number. The link state sequence number is used to detect old or duplicate LSAs.
- Checksum—*Display only.* Displays the checksum of the contents of the LSA.
- Tag—*Display only.* Displays the external route tag, a 32-bit field attached to each external route. This is not used by the OSPF protocol itself.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Type 7

Monitoring > Routing > Routing > OSPF LSAs > Type 7

Type 7 LSAs are NSSA AS-external routes that are flooded by the ASBR. They are similar to Type 5 LSAs, but unlike Type 5 LSAs, which are flooded into multiple areas, Type 7 LSAs are only flooded into NSSAs. Type 7 LSAs are converted to Type 5 LSAs by ABRs before being flooded into the area backbone.

Fields

- Process—*Display only.* Displays the OSPF process for the LSA.
- Area—*Display only.* Displays the OSPF area for the LSA.
- Network—*Display only.* Displays the address of the external network.
- Advertiser—*Display only.* Displays the router ID of the ASBR that sent the LSA.
- Age—*Display only.* Displays the age of the link state.
- Sequence #—*Display only.* Displays the link state sequence number. The link state sequence number is used to detect old or duplicate LSAs.
- Checksum—*Display only.* Displays the checksum of the contents of the LSA.
- Tag—*Display only.* Displays the external route tag, a 32-bit field attached to each external route. This is not used by the OSPF protocol itself.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

OSPF Neighbors

Monitoring > Routing > Routing > OSPF Neighbors

The OSPF Neighbor pane displays the OSPF neighbors dynamically discovered and statically configured OSPF neighbors on the security appliance.

Fields

- Neighbor—*Display only.* Displays the neighbor router ID.
- Priority—*Display only.* Displays the router priority.
- State—*Display only.* Displays the OSPF state for the neighbor:
  - Down—This is the first OSPF neighbor state. It means that no hello packets have been received from this neighbor, but hello packets can still be sent to the neighbor in this state.

    During the fully adjacent neighbor state, if the security appliance does not receive hello packet from a neighbor within the dead interval time, or if the manually configured neighbor is being removed from the configuration, then the neighbor state changes from Full to Down.
  - Attempt—This state is only valid for manually configured neighbors in an NBMA environment. In Attempt state, the security appliance sends unicast hello packets every poll interval to the neighbor from which hellos have not been received within the dead interval.

- Init—This state specifies that the security appliance has received a hello packet from its neighbor, but the ID of the receiving router was not included in the hello packet. When a router receives a hello packet from a neighbor, it should list the router ID of the sender in its hello packet as an acknowledgment that it received a valid hello packet.
- 2-Way—This state designates that bi-directional communication has been established between the security appliance and the neighbor. Bi-directional means that each device has seen the hello packet from the other device. This state is attained when the router receiving the hello packet sees its own Router ID within the neighbor field of the received hello packet. At this state, the security appliance decides whether to become adjacent with this neighbor. On broadcast media and non-broadcast multiaccess networks, a the security appliance becomes full only with the designated router and the backup designated router; it stays in the 2-way state with all other neighbors. On point-to-point and point-to-multipoint networks, the security appliance becomes full with all connected neighbors.

  At the end of this stage, the DR and BDR for broadcast and non-broadcast multiaccess networks are elected.

Note Receiving a Database Descriptor packet from a neighbor in the Init state will also a cause a transition to 2-way state.

- Exstart—Once the DR and BDR are elected, the actual process of exchanging link state information begins between the security appliance and the DR and BDR.

  In this state, the security appliance and the DR and BDR establish a master-slave relationship and choose the initial sequence number for adjacency formation. The device with the higher router ID becomes the master and starts the exchange and is therefore the only device that can increment the sequence number.

Note DR/BDR election occurs by virtue of a higher priority configured on the device instead of highest router ID. Therefore, it is possible that a DR plays the role of slave in this state. Master/slave election is on a per-neighbor basis. If multiple devices have the same DR priority, then the device with the highest IP address becomes the DR.

- Exchange—In the exchange state, OSPF neighbors exchange DBD packets. Database descriptors contain LSA headers only and describe the contents of the entire link state database. Each DBD packet has a sequence number which can be incremented only by master which is explicitly acknowledged by slave. Routers also send link state request packets and link state update packets (which contain the entire LSA) in this state. The contents of the DBD received are compared to the information contained in the routers link state database to check if new or more current link state information is available with the neighbor.
- Loading—In this state, the actual exchange of link state information occurs. Based on the information provided by the DBDs, routers send link state request packets. The neighbor then provides the requested link state information in link state update packets. During the adjacency, if a the security appliance receives an outdated or missing LSA, it requests that LSA by sending a link state request packet. All link state update packets are acknowledged.
- Full—In this state, the neighbors are fully adjacent with each other. All the router and network LSAs are exchanged and the router databases are fully synchronized.

  Full is the normal state for an OSPF router. The only exception to this is the 2-way state, which is normal in a broadcast network. Routers achieve the full state with their DR and BDR only. Neighbors always see each other as 2-way.

- Dead Time—*Display only.* Displays the amount of time remaining that the router waits to receive an OSPF hello packet from the neighbor before declaring the neighbor down.
- Address—*Display only.* Displays the IP address of the interface to which this neighbor is directly connected.
- Interface—*Display only.* Displays the interface on which the OSPF neighbor has formed adjacency.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Routes

Monitoring > Routing > Routing > Routes

The Routes pane displays the statically configured, connected, and discovered routes in the security appliance routing table.

Fields

- Protocol—*Display only.* Displays the origin of the route information.
  - RIP—The route was derived using RIP.
  - OSPF—The route was derived using OSPF.
  - CONNECTED—The route is a network directly connected to the interface.
  - STATIC—The route is statically defined.
- Type—*Display only.* Displays the type of route. It can be one of the following values:
  - - (dash)—Indicates that the type column does not apply to the specified route.
  - IA—The route is an OSPF interarea route.
  - E1—The route is an OSPF external type 1 route.
  - E2—The route is an OSPF external type 2 route.
  - N1—The route is an OSPF not so stubby area (NSSA) external type 1 route.
  - N2—The route is an OSPF NSSA external type 2 route.
- Destination—*Display only.* Displays the IP address/netmask of the destination network.
- Gateway—*Display only.* Displays the IP address of the next router to the remote network.
- Interface—*Display only.* Displays the interface through which the specified network can be reached.
- [AD/Metric]—*Display only.* Displays the administrative distance/metric for the route.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

CHAPTER 42

Monitoring VPN

The VPN Monitoring sections show parameters and statistics for the following:

- VPN statistics for specific Remote Access, LAN-to-LAN, WebVPN, and E-mail Proxy sessions
- Encryption statistics for tunnel groups
- Protocol statistics for tunnel groups
- Global IPSec and IKE statistics
- Crypto statistics for IPSec, IKE, SSL, and other protocols
- Statistics for cluster VPN server loads

VPN Connection Graphs

Displays VPN connection data in graphical or tabular form for the security appliance.

IPSec Tunnels

Monitoring > VPN > VPN Connection Graphs > IPSec Tunnels

Use this window to specify graphs and tables of the IPSec tunnel types you want to view, or prepare to export or print.

Fields

- Graph Window Title—Displays the default title that appears in the window when you click Show Graphs. This attribute is particularly useful when you want to clarify data in that window before printing or exporting it. To change the title, select an alternative from the drop-down list or type the title.
- Available Graphs—Shows the types of active tunnels you can view. For each type you want to view collectively in a single window, click the entry in this box and click Add.
- Selected Graphs—Shows the types of tunnels selected.

  If you click Show Graphs, ASDM shows the active tunnels types listed in this box in a single window.

  A highlighted entry indicates the type of tunnel to be removed from the list if you click Remove.
- Add—Moves the selected tunnel type from the Available Graphs box to the Selected Graphs box.

- Remove—Moves the selected tunnel type from the Selected Graphs box to the Available Graphs box.
- Show Graphs—Displays a window consisting of graphs of the tunnel types displayed in the Selected Graphs box. Each type in the window displayed has a Graph tab and a Table tab you can click to alternate the representation of active tunnel data.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Sessions

Monitoring > VPN > VPN Connection Graphs > Sessions

Use this panel to specify graphs and tables of the VPN session types you want to view, or prepare to export or print.

Fields

- Graph Window Title—Displays the default title that appears in the window when you click Show Graphs. This attribute is particularly useful when you want to clarify data in that window before printing or exporting it. To change the title, select an alternative from the drop-down list or type the title.
- Available Graphs—Shows the types of active sessions you can view. For each type you want to view collectively in a single window, click the entry in this box and click Add.
- Selected Graphs—Shows the types of active sessions selected.

  If you click Show Graphs, ASDM shows all of the active session types listed in this box in a single window.

  A highlighted entry indicates the type of session to be removed from the list if you click Remove.
- Add—Moves the selected session type from the Available Graphs box to the Selected Graphs box.
- Remove—Moves the selected session type from the Selected Graphs box to the Available Graphs box.
- Show Graphs—Displays a window consisting of graphs of the session types displayed in the Selected Graphs box. Each type in the window displayed has a Graph tab and a Table tab you can click to alternate the representation of active session data.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

VPN Statistics

These panels show detailed parameters and statistics for a specific remote-access, LAN-to-LAN, WebVPN, or E-mail Proxy session. The parameters and statistics differ depending on the session protocol. The contents of the statistical tables depend on the type of connection you select. The detail tables show all the relevant parameters for each session.

Sessions

Monitoring > VPN > VPN Statistics > Sessions

Use this panel to view session statistics for this server.

Fields

- Session types (unlabeled)—Lists the number of currently active sessions of each type, the total limit, and the total cumulative session count.
  - Remote Access—Shows the number of remote access sessions.
  - LAN-to-LAN—Shows the number of LAN-to-LAN sessions.
  - WebVPN—Shows the number of WebVPN sessions.
  - SSL VPN Client—Shows the number of SSL VPN Client (SVC) sessions.
  - E-mail Proxy—Shows the number of E-mail proxy sessions.
  - Total—Shows the total number of active concurrent sessions.
  - Total Cumulative—Shows the cumulative number of sessions since the last time the security appliance was rebooted or reset.
- Filter By—Specifies the type of sessions that the statistics in the following table represent.
  - Session type (unlabeled)—Designates the session type that you want to monitor. The default is Remote Access.
  - Session filter (unlabeled)—Designates which of the column heads in the following table to filter on. The default is --All Sessions--.
  - Filter name (unlabeled)—Specifies the name of the filter to apply. If you specify --All Sessions-- as the session filter list, this field is not available. For all other session filter selections, this field cannot be blank.
  - Filter—Executes the filtering operation.

The contents of the second table, also unlabeled, on this panel depend on the selection in the Filter By list. In the following list, the first-level bullets show the Filter By selection, and the second-level bullets show the column headings for this table.

- Remote Access—Indicates that the values in this table relate to remote access traffic.
  - Username/Tunnel Group—Shows the username or login name and the tunnel group for the session. If the client is using a digital certificate for authentication, the field shows the Subject CN or Subject OU from the certificate.
  - Assigned IP Address/Public IP Address—Shows the private (“assigned”) IP address assigned to the remote client for this session. This is also known as the “inner” or “virtual” IP address, and it lets the client appear to be a host on the private network. Also shows the Public IP address of the client for this remote-access session. This is also known as the “outer” IP address. It is typically assigned to the client by the ISP, and it lets the client function as a host on the public network.
  - Protocol/Encryption—Shows the protocol and the data encryption algorithm this session is using, if any.
  - Login Time/Duration—Shows the date and time (MMM DD HH:MM:SS) that the session logged in. and the length of the session. Time is displayed in 24-hour notation.
  - Client Type/Version—Shows the type and software version number (for example, rel. 7.0_int 50) for connected clients, sorted by username.
  - Bytes Tx/Bytes Rx—Shows the total number of bytes transmitted to/received from the remote peer or client by the security appliance.
  - NAC Result and Posture Token—Displays values in this column only if you configured Network Admission Control on the security appliance.

    The NAC Result shows one of the following values:

    Accepted—ACS successfully validated the posture of the remote host.

    Rejected—ACS could not successfully validate the posture of the remote host.

    Exempted—The remote host is exempt from posture validation according to the Posture Validation Exception list configured on the security appliance.

    Non-Responsive—The remote host did not respond to the EAPoUDP Hello message.

    Hold-off—The security appliance lost EAPoUDP communication with the remote host after successful posture validation.

    N/A—NAC is disabled for the remote host according to the VPN NAC group policy.

    Unknown—Posture validation is in progress.

    The posture token is an informational text string that is configurable on the Access Control Server. ACS downloads the posture token to the security appliance for informational purposes to aid in system monitoring, reporting, debugging, and logging. The typical value of the Posture Token field that follows the NAC Result field is as follows: Healthy, Checkup, Quarantine, Infected, or Unknown.
- LAN-to-LAN—Indicates that the values in this table relate to LAN-to-LAN traffic.
  - Tunnel Group/IP Address—Shows the name of the tunnel group and the IP address of the peer.
  - Protocol/Encryption—Shows the protocol and the data encryption algorithm this session is using, if any.
  - Login Time/Duration—Shows the date and time (MMM DD HH:MM:SS) that the session logged in. and the length of the session. Time is displayed in 24-hour notation.

  - Bytes Tx/Bytes Rx—Shows the total number of bytes transmitted to/received from the remote peer or client by the security appliance.
- WebVPN—Indicates that the values in this table relate to WebVPN traffic.
  - Username/IP Address—Shows the username or login name for the session and the IP address of the client.
  - Protocol/Encryption—Shows the protocol and the data encryption algorithm this session is using, if any.
  - Login Time/Duration—Shows the date and time (MMM DD HH:MM:SS) that the session logged in. and the length of the session. Time is displayed in 24-hour notation.
  - Client Type/Version—Shows the type and software version number (for example, rel. 7.0_int 50) for connected clients, sorted by username.
  - Bytes Tx/Bytes Rx—Shows the total number of bytes transmitted to/received from the remote peer or client by the security appliance.
- E-Mail Proxy—Indicates that the values in this table relate to WebVPN traffic.
  - Username/IP Address—Shows the username or login name for the session and the IP address of the client.
  - Protocol/Encryption—Shows the protocol and the data encryption algorithm this session is using, if any.
  - Login Time/Duration—Shows the date and time (MMM DD HH:MM:SS) that the session logged in. and the length of the session. Time is displayed in 24-hour notation.
  - Client Type/Version—Shows the type and software version number (for example, rel. 7.0_int 50) for connected clients, sorted by username.
  - Bytes Tx/Bytes Rx—Shows the total number of bytes transmitted to/received from the remote peer or client by the security appliance.

The remainder of this section describes the buttons and fields beside and below the table.

- Details—Displays the details for the selected session. The parameters and values differ, depending on the type of session.
- Logout—Ends the selected session.
- Ping—Sends an ICMP `ping` (Packet Internet Groper) packet to test network connectivity. Specifically, the security appliance sends an ICMP Echo Request message to a selected host. If the host is reachable, it returns an Echo Reply message, and the security appliance displays a Success message with the name of the tested host, as well as the elapsed time between when the request was sent and the response received. If the system is unreachable for any reason, (for example: host down, ICMP not running on host, route not configured, intermediate router down, or network down or congested), the security appliance displays an Error screen with the name of the tested host.
- Logout By—Selects a criterion to use to filter the sessions to be logged out. If you select any but --All Sessions--, the box to the right of the Logout By list becomes active. If you selected the value Protocol for Logout By, the box becomes a list, from which you can select a protocol type to use as the logout filter. The default value of this list is IPSec. For all choices other than Protocol, you must supply an appropriate value in this box.
- Logout Sessions—Ends all sessions that meet the specified Logout By criteria.
- Refresh—Updates the screen and its data. The date and time indicate when the screen was last updated.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Sessions Details

Monitoring > VPN > VPN Statistics > Sessions >Details

The Session Details window displays configuration settings, statistics, and state information about the selected session.

The Remote Detailed table at the top of the Session Details window displays the following columns:

- Username—Shows the username or login name associated with the session. If the remote peer is using a digital certificate for authentication, the field shows the Subject CN or Subject OU from the certificate.
- Group Policy and Tunnel Group—Group policy assigned to the session and the name of the tunnel group upon which the session is established.
- Assigned IP Address and Public IP Address—Private IP address assigned to the remote peer for this session. Also called the inner or virtual IP address, the assigned IP address lets the remote peer appear to be on the private network. The second field shows the public IP address of the remote computer for this session. Also called the outer IP address, the public IP address is typically assigned to the remote computer by the ISP. It lets the remote computer function as a host on the public network.
- Protocol/Encryption—Protocol and the data encryption algorithm this session is using, if any.
- Login Time and Duration—Time and date of the session initialization, and the length of the session. The session initialization time is in 24-hour notation.
- Client Type and Version—Type and software version number (for example, rel. 7.0_int 50) of the client on the remote computer.
- Bytes Tx and Bytes Rx—Shows the total number of bytes transmitted to and received from the remote peer by the security appliance.
- NAC Result and Posture Token—The ASDM displays values in this column only if you configured Network Admission Control on the security appliance.

  The NAC Result shows one of the following values:

  - Accepted—The ACS successfully validated the posture of the remote host.
  - Rejected—The ACS could not successfully validate the posture of the remote host.
  - Exempted—The remote host is exempt from posture validation according to the Posture Validation Exception list configured on the security appliance.
  - Non-Responsive—The remote host did not respond to the EAPoUDP Hello message.
  - Hold-off—The security appliance lost EAPoUDP communication with the remote host after successful posture validation.

  - N/A—NAC is disabled for the remote host according to the VPN NAC group policy.
  - Unknown—Posture validation is in progress.

  The posture token is an informational text string which is configurable on the Access Control Server. The ACS downloads the posture token to the security appliance for informational purposes to aid in system monitoring, reporting, debugging, and logging. The typical posture token that follows the NAC result is as follows: Healthy, Checkup, Quarantine, Infected, or Unknown.

The Details tab in the Session Details window displays the following columns:

- ID—Unique ID dynamically assigned to the session. The ID serves as the security appliance index to the session. It uses this index to maintain and display information about the session.
- Type—Type of session: IKE, IPSec, or NAC.
- Local Addr., Subnet Mask, Protocol, Port, Remote Addr., Subnet Mask, Protocol, and Port—Addresses and ports assigned to both the actual (Local) peer and those assigned to this peer for the purpose of external routing.
- Encryption—Data encryption algorithm this session is using, if any.
- Assigned IP Address and Public IP Address—Shows the private IP address assigned to the remote peer for this session. Also called the inner or virtual IP address, the assigned IP address lets the remote peer appear to be on the private network. The second field shows the public IP address of the remote computer for this session. Also called the outer IP address, the public IP address is typically assigned to the remote computer by the ISP. It lets the remote computer function as a host on the public network.
- Other—Miscellaneous attributes associated with the session.

  The following attributes apply to an IKE session:

  The following attributes apply to an IPSec session:

  The following attributes apply to a NAC session:

  - Revalidation Time Interval— Interval in seconds required between each successful posture validation.
  - Time Until Next Revalidation—0 if the last posture validation attempt was unsuccessful. Otherwise, the difference between the Revalidation Time Interval and the number of seconds since the last successful posture validation.
  - Status Query Time Interval—Time in seconds allowed between each successful posture validation or status query response and the next status query response. A status query is a request made by the security appliance to the remote host to indicate whether the host has experienced any changes in posture since the last posture validation.
  - EAPoUDP Session Age—Number of seconds since the last successful posture validation.
  - Hold-Off Time Remaining—0 seconds if the last posture validation was successful. Otherwise, the number of seconds remaining before the next posture validation attempt.
  - Posture Token—Informational text string configurable on the Access Control Server. The ACS downloads the posture token to the security appliance for informational purposes to aid in system monitoring, reporting, debugging, and logging. A typical posture token is Healthy, Checkup, Quarantine, Infected, or Unknown.
  - Redirect URL—Following posture validation or clientless authentication, the ACS downloads the access policy for the session to the security appliance. The Redirect URL is an optional part of the access policy payload. The security appliance redirects all HTTP (port 80) and HTTPS

(port 443) requests for the remote host to the Redirect URL if it is present. If the access policy does not contain a Redirect URL, the security appliance does not redirect HTTP and HTTPS requests from the remote host.

Redirect URLs remain in force until either the IPSec session ends or until posture revalidation, for which the ACS downloads a new access policy that can contain a different redirect URL or no redirect URL.

More—Press this button to revalidate or initialize the session or tunnel group.

The ACL tab displays the ACL containing the ACEs that matched the session.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Sub-session Details – NAC Details

Monitoring > VPN > VPN Statistics > Sessions >Details >More

The NAC Details window lets you view the statistics and state of a NAC session, and revalidate and initialize the session or tunnel group.

The statistics and state attributes in this window are as follows:

- Reval Int (T)—Revalidation Time Interval. Interval in seconds required between each successful posture validation.
- Reval Left (T)—Time Until Next Revalidation. 0 if the last posture validation attempt was unsuccessful. Otherwise, the difference between the Revalidation Time Interval and the number of seconds since the last successful posture validation.
- SQ Int (T)—Status Query Time Interval. Time in seconds allowed between each successful posture validation or status query response and the next status query response. A status query is a request made by the security appliance to the remote host to indicate whether the host has experienced any changes in posture since the last posture validation.
- EoU Age (T)—EAPoUDP Session Age. Number of seconds since the last successful posture validation.
- Hold Left (T)—Hold-Off Time Remaining. 0 seconds if the last posture validation was successful. Otherwise, the number of seconds remaining before the next posture validation attempt.
- Posture Token—Informational text string configurable on the Access Control Server. The ACS downloads the posture token to the security appliance for informational purposes to aid in system monitoring, reporting, debugging, and logging. A typical posture token is Healthy, Checkup, Quarantine, Infected, or Unknown.
- Redirect URL—Following posture validation or clientless authentication, the ACS downloads the access policy for the session to the security appliance. The Redirect URL is an optional part of the access policy payload. The security appliance redirects all HTTP (port 80) and HTTPS (port 443)

requests for the remote host to the Redirect URL if it is present. If the access policy does not contain a Redirect URL, the security appliance does not redirect HTTP and HTTPS requests from the remote host.

Redirect URLs remain in force until either the IPSec session ends or until posture revalidation, for which the ACS downloads a new access policy that can contain a different redirect URL or no redirect URL.

The buttons in this window are as follows:

Note Choose **Monitoring > VPN > VPN Statistics > NAC Session Summary** if you want to revalidate or initialize all sessions that are subject to posture validation.

- Revalidate Session—Click if the posture of the peer or the assigned access policy (that is, the downloaded ACL, if any) has changed. Clicking this button initiates a new, unconditional posture validation. The posture validation and assigned access policy that were in effect before you clicked this button remain in effect until the new posture validation succeeds or fails. Clicking this button does not affect the session if it is exempt from posture validation.
- Initialize Session—Click if the posture of the peer or the assigned access policy (that is, the downloaded ACL, if any) has changed, and you want to clear the resources assigned to the session. Clicking the button purges the EAPoUDP association and access policy, and initiates a new, unconditional posture validation. The NAC default ACL is effective during the revalidation, so the session initialization can disrupt user traffic. Clicking this button does not affect the session if it is exempt from posture validation.
- Revalidate Tunnel Group—Click if the posture of the peers in the tunnel group occupied by the selected session or the assigned access policies (that is, the downloaded ACLs), have changed. Clicking this button initiates new, unconditional posture validations. The posture validation and assigned access policy that were in effect for each session in the tunnel group before you clicked this button remain in effect until the new posture validation succeeds or fails. Clicking this button does not affect sessions that are exempt from posture validation.
- Initialize Tunnel Group—Click if the posture of the peers in the tunnel group occupied by the selected session, or the assigned access policies (that is, the downloaded ACLs), have changed, and you want to clear the resources assigned to the sessions. Clicking this button purges the EAPoUDP associations and access policies (that is, the downloaded ACLs, if any) used for posture validation in the tunnel group occupied by the selected session, and initiates new, unconditional posture validations for the effected peers. The NAC default ACL is effective during the revalidations, so the session initializations can disrupt user traffic. Clicking this button does not affect sessions that are exempt from posture validation.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Encryption Statistics

Monitoring > VPN > VPN Statistics > Encryption Statistics

This panel shows the data encryption algorithms used by currently active user and administrator sessions on the security appliance. Each row in the table represents one encryption algorithm type.

Fields

- Show Statistics For—Selects a specific server or group or all tunnel groups.
- Encryption Statistics—Shows the statistics for all the data encryption algorithms in use by currently active sessions.
  - Encryption Algorithm—Lists the encryption algorithm to which the statistics in this row apply.
  - Sessions—Lists the number of sessions using this algorithm.
  - Percentage—Indicates the percentage of sessions using this algorithm relative to the total active sessions, as a number. The sum of this column equals 100 percent (rounded).
- Total Active Sessions—Shows the number of currently active sessions.
- Cumulative Sessions—Shows the total number of sessions since the security appliance was last booted or reset.
- Refresh—Updates the statistics shown in the Encryption Statistics table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

NAC Session Summary

Monitoring > VPN > VPN Statistics > NAC Session Summary

The NAC Session Summary window lets you view the active and cumulative Network Admission Control sessions.

Fields

- Active NAC Sessions—General statistics about remote peers that are subject to posture validation.
- Cumulative NAC Sessions—General statistics about remote peers that are or have been subject to posture validation.
- Accepted—Number of peers that passed posture validation and have been granted an access policy by an Access Control Server.
- Rejected—Number of peers that failed posture validation or were not granted an access policy by an Access Control Server.
- Exempted—Number of peers that are not subject to posture validation because they match an entry in the Posture Validation Exception list configured on the security appliance.

- Non-responsive—Number of peers not responsive to Extensible Authentication Protocol (EAP) over UDP requests for posture validation. Peers on which no CTA is running do not respond to these requests. If the security appliance configuration supports clientless hosts, the Access Control Server downloads the access policy associated with clientless hosts to the security appliance for these peers. Otherwise, the security appliance assigns the NAC default policy.
- Hold-off—Number of peers for which the security appliance lost EAPoUDP communications after a successful posture validation. The NAC Hold Timer attribute (Configuration > VPN > NAC) determines the delay between this type of event and the next posture validation attempt.
- N/A—Number of peers for which NAC is disabled according to the VPN NAC group policy.
- Revalidate All—Click if the posture of the peers or the assigned access policies (that is, the downloaded ACLs), have changed. Clicking this button initiates new, unconditional posture validations of all NAC sessions managed by the security appliance. The posture validation and assigned access policy that were in effect for each session before you clicked this button remain in effect until the new posture validation succeeds or fails. Clicking this button does not affect sessions that are exempt from posture validation.
- Initialize All—Click if the posture of the peers or the assigned access policies (that is, the downloaded ACLs) have changed, and you want to clear the resources assigned to the sessions. Clicking this button purges the EAPoUDP associations and assigned access policies used for posture validations of all NAC sessions managed by the security appliance, and initiates new, unconditional posture validations. The NAC default ACL is effective during the revalidations, so the session initializations can disrupt user traffic. Clicking this button does not affect sessions that are exempt from posture validation.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Protocol Statistics

Monitoring > VPN > VPN Statistics > Protocol Statistics

This panel displays the protocols used by currently active user and administrator sessions on the security appliance. Each row in the table represents one protocol type.

Fields

- Show Statistics For—Selects a specific server or group or all tunnel groups.
- Protocol Statistics—Shows the statistics for all the protocols in use by currently active sessions.
  - Protocol—Lists the protocol to which the statistics in this row apply.
  - Sessions—Lists the number of sessions using this protocol.
  - Percentage—Indicates the percentage of sessions using this protocol relative to the total active sessions, as a number. The sum of this column equals 100 percent (rounded).
- Total Active Sessions—Shows the number of currently active sessions.

- Cumulative Sessions—Shows the total number of sessions since the security appliance was last booted or reset.
- Refresh—Updates the statistics shown in the Protocol Statistics table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Global IKE/IPSec Statistics

Monitoring > VPN > VPN Statistics > Global IKE/IPSec Statistics

This panel displays the global IKE/IPSec statistics for currently active user and administrator sessions on the security appliance. Each row in the table represents one global statistic.

Fields

- Show Statistics For—Selects a specific protocol, IKE Protocol (the default) or IPSec Protocol.
- Global IKE/IPSec Statistics—Shows the statistics for all the protocols in use by currently active sessions.
  - Statistic—Lists the name of the statistical variable. The contents of this column vary, depending upon the value you select for the Show Statistics For parameter.
  - Value—The numerical value for the statistic in this row.
- Refresh—Updates the statistics shown in the Global IKE/IPSec Statistics table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Crypto Statistics

Monitoring > VPN > VPN Statistics > Crypto Statistics

This panel displays the crypto statistics for currently active user and administrator sessions on the security appliance. Each row in the table represents one crypto statistic.

Fields

- Show Statistics For—Selects a specific protocol, IKE Protocol (the default), IPSec Protocol, SSL Protocol, or other protocols.
- Crypto Statistics—Shows the statistics for all the protocols in use by currently active sessions.
  - Statistic—Lists the name of the statistical variable. The contents of this column vary, depending upon the value you select for the Show Statistics For parameter.
  - Value—The numerical value for the statistic in this row.
- Refresh—Updates the statistics shown in the Crypto Statistics table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Compression Statistics

Monitoring > VPN > VPN Statistics > Compression Statistics

This panel displays the compression statistics for currently active user and administrator sessions on the security appliance. Each row in the table represents one compression statistic.

Fields

- Show Statistics For—Lets you select compression statistics for all VPN types, or for a single VPN type including IPSec VPN, WebVPN, or SSL VPN Client.
- Statistics—Shows all the statistics for the selected VPN type.
  - Statistic—Lists the name of the statistical variable. The contents of this column vary, depending upon the value you select for the Show Statistics For parameter.
  - Value—The numerical value for the statistic in this row.
- Refresh—Updates the statistics shown in the Compression Statistics table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Cluster Loads

Monitoring > VPN > VPN Statistics > Cluster Loads

Use this panel to view the current traffic load distribution among the servers in a VPN load-balancing cluster. If the server is not part of a cluster, you receive an information message saying that this server does not participate in a VPN load-balancing cluster.

Fields

- VPN Cluster Loads—Displays the current load distribution in the VPN load-balancing cluster. Clicking a column heading sorts the table, using the selected column as the sort key.
  - Public IP Address—Displays the externally visible IP address for the server.
  - Role—Indicates whether this server is a master or backup device in the cluster.
  - Priority—Shows the priority assigned to this server in the cluster. The priority must be an integer in the range of 1 (lowest) to 10 (highest). The priority is used in the master-election process as one way to determine which of the devices in a VPN load-balancing cluster becomes the master or primary device for the cluster.
  - Model—Indicates the security appliance model name and number for this server.
  - Load %—Indicates what percentage of a server's total capacity is in use, based upon the capacity of that server.
  - Sessions—Shows the number of currently active sessions.
- Refresh—Loads the table with updated statistics.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

WebVPN SSO Statistics

Monitoring > VPN > WebVPN > SSO Statistics

This panel displays the single sign-on statistics for currently active SSO servers configured for the security appliance.

Note These statistics are for SSO with SiteMinder servers only.

Fields

- Show Statistics For—Selects an SSO server.
- SSO Statistics—Shows the statistics for all the currently active sessions on the selected SSO server.

  SSO statistics that display include:

  - Name of SSO server
  - Type of SSO server
  - Authentication Scheme Version
  - Web Agent URL
  - Number of pending requests
  - Number of authorization requests
  - Number of retransmissions
  - Number of accepts
  - Number of rejects
  - Number of timeouts
  - Number of unrecognized responses
- Refresh—Updates the statistics shown in the SSO Statistics table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

Monitoring Properties

Properties contains the following topics:

AAA Servers

Monitoring > Properties > AAA Servers

The AAA Server pane lets you view the AAA Server configuration.

Prerequisites

None.

Fields

The AAA Server pane displays the following fields:

- Server Group—Displays a configured server group, or LOCAL if none have been configured.
- Protocol—Displays what protocol the server group uses for AAA.
- IP Address—Displays the IP address of the configured AAA server.

Below the list of AAA servers are the statistics for each configured server. You can clear the statistics using the Clear Server Stats button.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

CRL

Monitoring > Properties > CRL

This pane allows you to view or clear associated CRLs of selected Trustpoints. Trustpoints are configured in Configuration > Device Administration > Certificates > Trustpoints.

Fields

- Trustpoint name—The name of the selected Trustpoint.
- View CRL—View the selected CRL.
- Clear CRL—Clear the selected CRL from the cache.
- CRL info—*Display only.* Displays detailed CRL information.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Connection Graphs

The Connection Graphs pane let you view connection information about the security appliance in graph format. You can view information about NAT and performance monitoring information, including UDP connections, AAA performance, and inspection information. Refer to the following topics for more information:

- Xlates
- Perfmon

Xlates

Monitoring > Properties > Connection Graphs > Xlates

Xlates lets you view the active Network Address Translations in a graph format. You can graph a maximum of four graphs in one frame.

Fields

- Available Graphs For:—Lists the components you can graph.
  - Xlate Utilization—Displays the security appliance NAT utilization.
- Graph Window—Shows the graph window name to which you want to add a statistic type. If you have a graph window already open, a new graph window is listed by default. If you want to add a statistic type to an already open graph, select the open graph window name. The statistics already included in the graph window are shown in the Selected Graphs field, to which you can add additional types (up to a maximum of four types per window).
- Add—Click this button to move the selected entries in the Available Graphs For field to the Selected Graphs field.
- Remove—Removes the selected statistic type from the Selected Graphs field.
- Show Graphs—Click to display a new or updated graph window with the selected statistics.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Perfmon

Monitoring > Properties > Connection Graphs > Perfmon

The Perfmon pane lets you view the performance information in a graph format. You can graph a maximum or four graphs in one frame.

This information includes the number of translations, connections, Websense requests, address translations, and AAA transactions that occur each second.

Fields

- Available Graphs For:—Lists the components you can graph.
  - AAA Perfmon—Displays the security appliance AAA performance information.
  - Inspection Perfmon—Displays the security appliance inspection performance information.
  - Web Perfmon—Displays the security appliance web performance information, including URL access and URL server requests.
  - Connections Perfmon—Displays the security appliance connections performance information.
  - Xlate Perfmon—Displays the security appliance NAT performance information.
- Graph Window—Shows the graph window name to which you want to add a statistic type. If you have a graph window already open, a new graph window is listed by default. If you want to add a statistic type to an already open graph, select the open graph window name. The statistics already included in the graph window are shown in the Selected Graphs field, to which you can add additional types (up to a maximum of four types per window).

- Add—Click this button to move the selected entries in the Available Graphs For field to the Selected Graphs field.
- Remove—Removes the selected statistic type from the Selected Graphs field.
- Show Graphs—Click to display a new or updated graph window with the selected statistics.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

DNS Cache

Monitoring > Properties > DNS Cache

The security appliance provides a local cache of DNS information from external DNS queries sent out for certain WebVPN and Certificate commands. Each DNS translation request is first looked for in the local cache. If the local cache has the information, the resulting IP address is returned. If the local cache can not resolve the request, a DNS query is sent to the various DNS servers that have been configured. If an external DNS server resolves the request, the resulting IP address is stored in the local cache along with its corresponding hostname.

Important Notes

- DNS cache entries are time stamped. The time stamp will be used to age out unused entries. When the entry is added to the cache the time stamp is initialized. Each time the entry is accessed the timestamp is updated. At a configured time interval, the DNS cache will check all entries and purge those entries whose time exceeds a configured age out timer.
- If new entries arrive but there is no room in the cache, since the size exceeded or there is no more memory allowed, the cache will be thinned by one third based on entries age. The oldest entries will be removed.
- The entire cache can be cleared by clicking the **Clear Cache** button.

Fields

- Host—The DNS name of the host.
- IP Address—Shows the address that resolves to the hostname.
- Permanent—Indicates if the entry made though a **name** command.
- Idle Time—Gives the time elapsed since the security appliance last referred to that entry.
- Active—Indicates if the entry has aged out. If there is no sufficient space in cache, this entry may be deleted.
- Clear Cache—Clears the DNS cache.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Device Access

Monitoring > Properties >Device Access

Device Access lets you monitor management sessions, AAA locked out users, and authenticated users.

Device Access contains the following topics:

- AAA Local Locked Out Users
- Authenticated Users
- HTTPS/ASDM Sessions
- Secure Shell Sessions
- Telnet Sessions

AAA Local Locked Out Users

Monitoring > Properties > Device Access > AAA Local Locked Out Users

The AAA Local Locked Out Users pane lets you view a list of users who have been locked out of ASDM for failed login attempts. You can also clear selected lockout conditions or all lockouts.

Fields

The AAA Local Lockouts area displays the following.

- Currently locked out users—A list of the currently locked out users.
- Lock Time—The amount of time the user has been locked out from accessing the system.
- Failed Attempts—The number of failed login attempts.
- User—The user name used with the failed login attempts.

The following buttons are also available:

- Refresh—Updates the display with the most current information.
- Clear lockout—Clears the selected user lockout condition.
- Clear all lockouts—Clears all user lockout conditions. It is good practice to refresh the list of lockout conditions before clearing all lockouts.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Authenticated Users

Monitoring > Properties > Device Access > Authenticated Users

This pane lets you view what users are authenticated to the security appliance.

Fields

- User—Displays the user name of the person authenticated to the security appliance.
- IP Address—Displays the IP address of the user authenticated to the security appliance.
- Dynamic ACL—Displays the dynamic access list of the user authenticated to the security appliance.
- Inactivity Timeout—Displays the amount of time the selected user must remain inactive before the session times out and the user is disconnected.
- Absolute Timeout—Displays the amount of time the selected user can remain connected before the session closes and the user is disconnected.
- Refresh—Select to refresh the list of currently authenticated users.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

HTTPS/ASDM Sessions

Monitoring > Properties > Device Access > HTTPS/ASDM

The HTTPS/ASDM pane lets you view currently connected HTTPS/ASDM sessions.

A secure connection is needed so that a PC or workstation client running ASDM in a network browser window can communicate with the security appliance.

Fields

The HTTPS/ASDM pane displays the following fields:

- Session ID—Displays the name of a connected HTTPS/ASDM session.
- IP Address—Displays the IP address of each host or network permitted to connect to this security appliance.

- Refresh—Select to refresh the list of currently connected HTTPS/ASDM sessions.
- Disconnect—Select to disconnect a connected HTTPS/ASDM session.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Secure Shell Sessions

Monitoring > Properties > Device Access > Secure Shell Sessions

The Secure Shell Sessions pane lets you view hosts connected to the security appliance for administrative access using the SSH protocol.

Fields

The Currently Connected Secure Shell Sessions pane displays the following fields:

- Client—Displays the client type for the selected SSH session.
- User—Displays the user name for the selected SSH session.
- State—Displays the state of the selected SSH session.
- Version—Displays the version of SSH used to connect to the security appliance.
- Encryption (in)—Displays the inbound encryption method used for the selected session.
- Encryption (out)—Displays the outbound encryption method used for the selected session.
- HMAC (in)—Displays the configured HMAC for the selected inbound SSH session.
- HMAC (out)—Displays the configured HMAC for the selected outbound SSH session.
- SID—Displays the secure ID of the selected session.
- Refresh—Select to refresh the list of currently connected SSH sessions.
- Disconnect—Select to disconnect a connected SSH session.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Telnet Sessions

Monitoring > Properties > Device Access > Telnet

The Telnet Sessions pane lets you view currently connected Telnet sessions.

Fields

The Telnet Sessions pane displays the following fields:

- Session ID—Displays the name of a connected Telnet sessions.
- IP Address—Displays the IP address of each host permitted to connect to this security appliance over Telnet.
- Refresh—Select to refresh the list of currently connected Telnet sessions.
- Disconnect—Select to disconnect a connected Telnet session.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

IP Audit

Monitoring > Properties > IP Audit

The IP Audit pane lets you view the number of packets that match informational and attack signatures in graph or table form. Each statistic type shows the combined packets for all interfaces that have IP audit enabled.

Fields

- Available Graphs for—Lists the types of signatures available for monitoring. See IP Audit Signatures fro detailed information about each signature type. 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.
  - IP Options—Shows the packet count for the following signatures:

    Bad Options List (1000)

    Timestamp (1002)

    Provide s, c, h, tcc (1003)

    SATNET ID (1005)
  - IP Route Options—Shows the packet count for the following signatures:

    Loose Source Route (1004)

    Record Packet Route (1001)

    Strict Source Route (1006)

- IP Attacks—Shows the packet count for the following signatures:
  IP Fragment Attack (1100)
  Impossible IP Packet (1102)
  IP Teardrop (1103)
- ICMP Requests—Shows the packet count for the following signatures:
  Echo Request (2004)
  Time Request (2007)
  Info Request (2009)
  Address Mask Request (2011)
- ICMP Responses—Shows the packet count for the following signatures:
  Echo Reply (2000)
  Source Quench (2002)
  Redirect (2003)
  Time Exceeded (2005)
  Parameter Problem (2006)
- ICMP Replies—Shows the packet count for the following signatures:
  Unreachable (2001)
  Time Reply (2008)
  Info Reply (2010)
  Address Mask reply (2012)
- ICMP Attacks—Shows the packet count for the following signatures:
  Fragmented ICMP (2150)
  Large ICMP (2151)
  Ping of Death (2154)
- TCP Attacks—Shows the packet count for the following signatures:
  No Flags (3040)
  SYN & FIN Flags Only (3041)
  FIN Flag Only (3042)
- UDP Attacks—Shows the packet count for the following signatures:
  Bomb (4050)
  Snork (4051)
  Chargen (4052)
- DNS Attacks—Shows the packet count for the following signatures:
  Host Info (6050)
  Zone Transfer (6051)
  Zone Transfer High Port (6052)
  All Records (6053)
- FTP Attacks—Shows the packet count for the following signatures:

Improper Address (3153)

Improper Port (3154)

- RPC Requests to Target Hosts—Shows the packet count for the following signatures:

  Port Registration (6100)

  Port Unregistration (6101)

  Dump (6102)

- YP Daemon Portmap Requests—Shows the packet count for the following signatures:

  ypserv Portmap Request (6150)

  ypbind Portmap Request (6151)

  yppasswdd Portmap Request (6152)

  ypupdated Portmap Request (6153)

  ypxfrd Portmap Request (6154)

- Miscellaneous Portmap Requests—Shows the packet count for the following signatures:

  mountd Portmap Request (6155)

  rexd Portmap Request (6175)

- Miscellaneous RPC Calls—Shows the packet count for the following signatures:

  rexd Attempt (6180)

- RPC Attacks—Shows the packet count for the following signatures:

  statd Buffer Overflow (6190)

  Proxied RPC (6103)

- Add—Adds the selected statistic type to the selected graph window.
- Remove—Removes the selected statistic type from the selected graph window.
- Show Graphs—Shows the graph window name to which you want to add a statistic type. If you have a graph window already open, a new graph window is listed by default. If you want to add a statistic type to an already open graph, select the open graph window name. The statistics already included on the graph are shown in the Selected Graphs pane, to which you can add additional types. Graph windows are named for ASDM followed by the interface IP address and the name “Graph”. Subsequent graphs are named “Graph (2)” and so on.
- Selected Graphs—Shows the statistic types you want to show in the selected graph window. You an include up to four types.
  - Show Graphs—Shows the graph window or updates the graph with additional statistic types if added.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

System Resources Graphs

Monitoring > Properties > System Resources Graphs

System Resources Graphs lets you view the status of the security appliance memory, CPU, and block utilization. You can graph a maximum or four graphs in one frame.

System Resources Graphs contains the following topics:

Blocks

Monitoring > Properties > System Resources Graphs > Blocks

Blocks lets you view the free and used memory blocks in a graph format. You can graph a maximum or four graphs in one frame.

Fields

- Available Graphs For:—Lists the components you can graph.
  - Blocks Used—Displays the security appliance used memory blocks.
  - Blocks Free—Displays the security appliance free memory blocks.
- Graph Window—Shows the graph window name to which you want to add a statistic type. If you have a graph window already open, a new graph window is listed by default. If you want to add a statistic type to an already open graph, select the open graph window name. The statistics already included in the graph window are shown in the Selected Graphs field, to which you can add additional types (up to a maximum of four types per window).
- Add—Click this button to move the selected entries in the Available Graphs For field to the Selected Graphs field.
- Remove—Removes the selected statistic type from the Selected Graphs field.
- Show Graphs—Click to display a new or updated graph window with the selected statistics.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

CPU

Monitoring > Properties > System Resources Graphs > CPU

CPU lets you view the CPU utilization in a graph format. You can graph a maximum or four graphs in one frame.

Fields

- Available Graphs For:—Lists the components you can graph.
  - CPU Utilization—Displays the security appliance CPU utilization.
- Graph Window—Shows the graph window name to which you want to add a statistic type. If you have a graph window already open, a new graph window is listed by default. If you want to add a statistic type to an already open graph, select the open graph window name. The statistics already included in the graph window are shown in the Selected Graphs field, to which you can add additional types (up to a maximum of four types per window).
- Add—Click this button to move the selected entries in the Available Graphs For field to the Selected Graphs field.
- Remove—Removes the selected statistic type from the Selected Graphs field.
- Show Graphs—Click to display a new or updated graph window with the selected statistics.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Memory

Monitoring > Properties > System Resources Graphs > Memory

Memory lets you view the memory utilization in a graph format. You can monitor free and used memory available in real time. You can graph a maximum or four graphs in one frame.

Fields

- Available Graphs For:—Lists the components you can graph.
  - Free Memory—Displays the security appliance free memory.
  - Used Memory—Displays the security appliance used memory.
- Graph Window—Shows the graph window name to which you want to add a statistic type. If you have a graph window already open, a new graph window is listed by default. If you want to add a statistic type to an already open graph, select the open graph window name. The statistics already included in the graph window are shown in the Selected Graphs field, to which you can add additional types (up to a maximum of four types per window).
- Add—Click this button to move the selected entries in the Available Graphs For field to the Selected Graphs field.

- Remove—Removes the selected statistic type from the Selected Graphs field.
- Show Graphs—Click to display a new or updated graph window with the selected statistics.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

A

B

C

D

E

F

G

H

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X

Z