Cisco Security Appliance Configuration Guide using ASDM

Version 6.2

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

Text Part Number: OL-18494-03

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

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

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

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

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

CCDE, CCENT, CCSI, Cisco Eos, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco:Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumin, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLYNX, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0910R)

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

Cisco Security Appliance Configuration Guide using ASDM

CONTENTS

Getting Started

CHAPTER 3

CHAPTER 4

Device Setup and Management

CHAPTER 7

CHAPTER 15

CHAPTER 16

Configuring the Firewall

Configuring VPN

Monitoring the Security Appliance

Reference

Preface

The *Cisco Security Appliance Configuration Guide using ASDM* contains the information that is available in the ASDM online help system.

This preface contains the following topics:

- Related Documentation, page xlv
- Document Conventions, page xlv
- Obtaining Documentation and Submitting a Service Request, page xlvi

Related Documentation

For more information, refer to the following documentation:

- *Release Notes for Cisco ASDM*
- *Cisco ASA 5500 Series Configuration Guide using the CLI*
- *Cisco ASA 5500 Series Command Reference*
- *Cisco ASA 5500 Series Adaptive Security Appliance Getting Started Guide*
- *Cisco ASA 5505 Series Adaptive Security Appliance Getting Started Guide*
- *Cisco ASA 5500 Series Release Notes*
- *Cisco ASA 5500 Series System Log Messages*
- *Open Source Software Licenses for ASA and PIX Security Appliances*

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 and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS Version 2.0.

PART 1

Getting Started

CHAPTER 1

Before You Start

This section describes the tasks you must perform before you use ASDM, and includes the following topics:

- Factory Default Configurations, page 1-1
- Configuring the Security Appliance for ASDM Access, page 1-4
- Setting Transparent or Routed Firewall Mode at the CLI, page 1-5
- Starting ASDM, page 1-5
- Configuration Overview, page 1-8

Factory Default Configurations

The factory default configuration is supported on all security appliances, except for the PIX 525 and PIX 535 models.

For the ASA 5505 model, the factory default configuration includes predefined interfaces and NAT, so that the adaptive security appliance is ready to use in your network as delivered.

For the PIX 515, PIX515E, ASA 5510, and higher version models, the factory default configuration provides a management interface to allow you to connect to the security appliance using ASDM, from which you can then complete your configuration.

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

This section includes the following topics:

- Restoring the Factory Default Configuration, page 1-1
- ASA 5505 Default Configuration, page 1-2
- ASA 5510 and Higher Version Default Configuration, page 1-3
- PIX 515/515E Default Configuration, page 1-4

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, do one of the following:

- For the ASA 5500 series, check the **Use this address for the Management 0/0 interface that will be named as "management"** check box, enter the new IP address in the Management IP Address field, and then choose the new subnet mask in the Management Subnet Mask drop-down list.
- For the PIX series, check the **Use this address for the Ethernet 1 interface, which will be named "inside"** check box, enter the new inside IP address in the Inside IP Address field, and then choose the new inside subnet mask in the Inside Subnet Mask drop-down list.

Step 3 Click **OK**.

Note After restoring the factory default configuration, the next time you reload the adaptive security appliance, it boots from the first image in internal Flash memory. If an image does not exist in internal Flash memory, the adaptive security appliance does not boot.

ASA 5505 Default Configuration

The default factory configuration for the ASA 5505 adaptive security appliance provides 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 interface using 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 adaptive security appliance, so that a computer connecting to the VLAN 1 interface receives an IP 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 Version Default Configuration

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

- The 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 adaptive security appliance, so a computer 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 provides 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 subnet mask are 192.168.1.1 and 255.255.255.0.
- The DHCP server is enabled on the security appliance, so a computer 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 and you have a factory default configuration, you can connect to the default management address by pointing your browser or the ASDM launcher to the IP address in the following URL:

https://192.168.1.1/admin

With the factory default configuration, clients on the 192.168.1.0/24 inside network can access ASDM. To allow other clients to access ASDM, see the "Configuring Device Access for ASDM, Telnet, or SSH" section on page 18-1.

See the following Ethernet connection guidelines when using the factory default configurations:

- ASA 5505—The switch port to which you connect to ASDM can be any port, except for Ethernet 0/0.
- ASA 5510 and higher —The interface to which you connect to ASDM is Management 0/0.
- PIX 515/515E—The interface to which you connect to ASDM is Ethernet 1.

Note Other PIX models do not have factory default configurations.

For more information, see the "Factory Default Configurations" section on page 1-1. If you do not have a factory default configuration, see the *Cisco ASA 5500 Series Configuration Guide using the CLI* for instructions to access the CLI and the **setup** command to perform minimum initial configuration.

Setting Transparent or Routed Firewall Mode at the CLI

By default, the security appliance runs in routed firewall mode. If you want to run in transparent firewall mode, you must set the mode using the CLI before you configure anything else on the security appliance. (When you change modes, the adaptive security appliance clears the configuration, because many commands are not supported in both modes.)

To set the mode at the CLI, see the

Starting ASDM

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

- Downloading the ASDM Launcher, page 1-5
- Starting ASDM from the ASDM Launcher, page 1-5
- Using ASDM in Demo Mode, page 1-6
- Starting ASDM from a Web Browser, page 1-7

Downloading the ASDM Launcher

The ASDM Launcher is for Windows only. The ASDM Launcher avoids double authentication and certificate dialog boxes, launches more quickly, and caches previously entered IP addresses and usernames.

To download the ASDM launcher, perform the following steps:

Step 1 On the ASDM Welcome screen, click the applicable button to download the ASDM Launcher installation file.

Step 2 Double-click the **asdm-launcher.exe** file.

Note In transparent firewall mode, enter the management IP address. Be sure to enter `https`, not `http`.

Step 3 Click **OK** or **Yes** to all prompts, including the name and password prompt. Leave the name and password blank.

The installer downloads to your computer.

Step 4 Run the installer to install the ASDM Launcher.

Starting ASDM from the ASDM Launcher

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

Step 1 Double-click the Cisco ASDM Launcher shortcut on your desktop, or open it from the **Start** menu. Alternatively, from the ASDM Welcome screen, you can click **Run Startup Wizard** to configure ASDM.

Step 2 Enter or choose the adaptive security appliance IP address or hostname to which you want to connect. To clear the list of IP addresses, click the trash can icon next to the Device/IP Address/Name field.

Step 3 Enter your username and your password, and then click **OK**.

If there is a new version of ASDM on the adaptive security appliance, the ASDM Launcher automatically downloads the new version and requests that you update the current version before starting ASDM.

Note If you are using the Factory Default Configuration, you do not need to have a username or password. Leave these fields blank to login to ASDM.

Using ASDM in Demo Mode

The ASDM Demo Mode, a separately installed application, lets you run ASDM without having a live device available. In this mode, you can do the following:

- Perform configuration and selected 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 CSC SSM.
- Obtain simulated monitoring and logging data, including real-time system log messages. The data shown is randomly generated; however, the experience is identical to what you would see when you are connected to a real device.

This mode does not support the following:

- Saving changes made to the configuration that appear in the GUI.
- File or disk operations.
- Historical monitoring data.
- Non-administrative users.
- These 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 Software

    File Transfer

    Upload image from Local PC

    System Reload

  - Toolbar/Status bar > Save
  - Configuration > Interface > Edit Interface > Renew DHCP Lease
  - Configuring a standby device after failover
- Operations that cause a rereading of the configuration, in which the GUI reverts to the original configuration:
  - Switching contexts
  - Making changes in the Interface pane
  - NAT pane changes
  - Clock pane changes

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

Step 1 Download the ASDM Demo Mode installer, asdm-demo-*version*.msi, from one of the following locations:

- http://www.cisco.com/cgi-bin/tablebuild.pl/asa
- http://www.cisco.com/cgi-bin/tablebuild.pl/pix

Step 2 Double-click the installer to install the software.

Step 3 Double-click the Cisco ASDM Launcher shortcut on your desktop, or open it from the **Start** menu.

Step 4 Check the **Run in Demo Mode** check box.

The Demo Mode window appears.

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`*

Where *interface_ip_address* is the IP address of ASDM on the adaptive security appliance network.

Note In transparent firewall mode, enter the management IP address. Be sure to enter **`https`**, not **`http`**.

Step 2 Click **OK** or **Yes** to all browser prompts, including the username and password, which you should leave blank.

The Cisco ASDM 6.2(5) Welcome page displays with the following buttons:

- **Install ASDM Launcher and Run ASDM**
- **Run ASDM**
- **Run Startup Wizard**

Step 3 Click **Run ASDM**.

Step 4 Click **OK** or **Yes** to all the browser prompts.

Configuration Overview

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

Step 1 For initial configuration Using the Startup Wizard, choose **Wizards** > **Startup Wizard**.

Step 2 To use the IPSec VPN Wizard to configure IPSec VPN connections, choose **Wizards** > **IPSec VPN Wizard** and complete each screen that appears.

Step 3 To use the SSL VPN Wizard to configure SSL VPN connections, choose **Wizards** > **SSL VPN Wizard** and complete each screen that appears.

Step 4 To configure high availability and scalability settings, choose **Wizards** > **High Availability and Scalability Wizard**. See Configuring Failover with the High Availability and Scalability Wizard for more information.

Step 5 To use the Configuring and Running Captures with the Packet Capture Wizard to configure packet capture, choose **Wizards** > **Packet Capture Wizard**.

Step 6 To display different colors and styles available in the ASDM GUI, choose **View > Office Look and Feel**.

Step 7 To configure features, click the **Configuration** button on the toolbar and then click one of the following feature buttons to display the associated configuration pane: **Device Setup**, **Device Management**, **Firewall**, **Remote Access VPN**, **Site-to-Site VPN, IPS,** and **Trend Micro Content Security**.

Note If the Configuration screen is blank, click **Refresh** on the toolbar to display the screen content.

- The Device Setup pane lets you do the following:
  - Launch the Startup Wizard to create security policy.
  - Configure basic interface parameters, including the IP address, name, security level, and the bridge group for transparent mode. For more information, see Configuring General Interface Parameters.
  - Configure OSPF, RIP, static, and asymmetric routing (single mode only). For more information, see Configuring Dynamic And Static Routing.
  - Configure AAA services.
  - Configure digital certificates.
  - Configure the device name and device password.
  - Configure DHCP services.
  - Configure DNS services.
- The Firewall pane lets you configure security policy, including access rules, AAA rules, filter rules, service policy rules, as well as NAT rules, URL filtering servers, global objects, and perform advanced configuration for the following:
  - Information About Access Rules and ACLs determine the access of IP traffic through the security appliance. For transparent firewall mode, you can also apply an EtherType access list to allow non-IP traffic.

  - Ethertype Rule Field Definitions (Transparent Mode Only) determine the access of non-IP traffic through the security appliance.
  - Configuring Access Rules and ACLs determine authentication and/or authorization for certain types of traffic, for example, HTTP. The security appliance also sends accounting information to a RADIUS or TACACS+ server.
  - Filter Rules prevent outbound access to specific websites or FTP servers. The security appliance works with a separate server running either Websense Enterprise or Sentian by N2H2. Choose **Configuration > Properties > URL Filtering** to configure the URL filtering server, which you must do before adding a rule.
  - Configuring Service Policy Rules apply 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 adaptive 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 a 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 setting lets you use private addresses, which are not routable on the Internet, on your inside networks.
  - Adding 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 adaptive security appliance. These reusable components, or objects, include the following:

    Network Objects/Groups

    Service Groups

    Class Maps

    Inspect Maps

    Regular Expressions

    TCP Maps

    Global Pools

    Time Ranges

- The Remote Access VPN pane lets you configure network client access, clientless SSL VPN browser access and advanced web-related settings, AAA setup, certificate management, load balancing, and perform additional advanced configuration, including the following:
  - Configure IPSec connections for VPN tunnels.
  - Configure clientless SSL VPN connections. Clientless SSL VPN lets users establish a secure, remote-access VPN tunnel to the adaptive security appliance using a web browser.
  - Configuring IKE, Load Balancing, and NAC sets the IP addresses of clients after they connect through the VPN tunnel.
  - Configuring Load Balancing configures load balancing for VPN connections.
  - E-Mail Proxy configures e-mail proxies. E-mail proxies extend remote e-mail capability to clientless SSL VPN users.
- The Site-to-Site VPN pane lets you configure site-to-site VPN connections, group policies, certificate management, and perform advanced configuration, including the following:

  - Configuring IKE, Load Balancing, and NAC Policies and IKE Parameters (also called ISAKMP), which provide the negotiation protocol that lets two hosts agree on how to build an IPSec security association.
- The Device Management pane lets you configure settings to access and manage the following:
  - ASDM and HTTP over SSL management sessions.
  - FTP and TFTP clients.
  - The CLI.
  - SNMP and ICMP.
  - Logging, including e-mail, event lists, filters, rate limit, syslog servers, and SMTP. For more information, see Configuring Logging.
  - User and AAA authentication.
  - High availability, the Scalability Wizard, and failover.
  - Advanced configuration.

Note If you have a CSC SSM card or IPS software installed, either the **Trend Micro Content Security** or **IPS** feature button also appears.

- The IPS pane lets you configure the IPS sensor. For more information, see Configuring the IPS Application on the AIP SSM and SSC.
- The Trend Micro Content Security pane lets you configure the CSC SSM (available for the ASA 5510, ASA 5520, and ASA 5540 adaptive security appliances). For more information, see Configuring Trend Micro Content Security.

Step 8 To monitor the adaptive security appliance, click the **Monitoring** button on the toolbar and then click one of the following feature buttons to display the associated monitoring pane: **Interfaces**, **VPN**, **Trend Micro Content Security**, **Routing**, **Properties**, and **Logging**.

- The Interfaces pane lets you monitor the ARP table, DHCP services, dynamic access lists, the PPoE client, connection status, and interface statistics. For more information, see Monitoring Interfaces.
- The VPN pane lets you monitor VPN connections. For more information, see Monitoring VPN.
- The Routing pane lets you Monitors routes, OSPF LSAs, and OSPF neighbors. For more information, see Monitoring Routing.
- The Properties pane lets you monitor management sessions, AAA servers, failover, CRLs, the DNS cache, and system statistics. For more information, see Monitoring Properties.
- The Logging pane lets you monitor system log messages, the Real-Time Log Viewer, and the log buffer. For more information, see Monitoring Logging.
- The Trend Micro Content Security pane lets you monitor CSC SSM connections. For more information, see Monitoring Trend Micro Content Security.

CHAPTER 2

Welcome to ASDM

Cisco Adaptive Security Device Manager (ASDM) delivers world-class security management and monitoring services for security appliances through an intuitive, easy-to-use, management interface. Bundled with supported security appliances, the device manager accelerates security appliance deployment with intelligent wizards, robust administration tools, and versatile monitoring services that complement the advanced security and networking features offered by Cisco ASA 5500 series and Cisco PIX 500 series security appliances.

Note If you change the color scheme of your operating system while ASDM is running, you should restart ASDM, because some ASDM windows might not display correctly.

This chapter includes the following sections:

ASDM Client Operating System and Browser Requirements

Table 2-1 lists the supported and recommended client operating systems and Java for ASDM.

Table 2-1 Operating System and Browser Requirements

	Browser			Sun Java SE Plug-in[1]
Operating System	Internet Explore	Firefox	Safari	
Microsoft Windows (English and Japanese): • 7 • Vista • 2003 Server • XP • 2000 (Service Pack 4 or higher)	6.0 or above	1.5 or above	No support.	• 5.0 (1.5.0) • 6.0
Apple Macintosh OS X: • 10.6 • 10.5 • 10.4	No support.	1.5 or above	2.0 or above	• 5.0 (1.5.0) • 6.0
Red Hat Enterprise Linux 5 (GNOME or KDE): • Desktop • WS	N/A	1.5 or above	N/A	• 5.0 (1.5.0) • 6.0

[1] Obtain Sun Java from java.sun.com.

Note ASDM supports up to a maximum of a 512 KB configuration. If you exceed this amount, you may experience performance issues.

Supported Platforms and SSMs

Note ASDM 6.2(1) and higher is not supported on the PIX platforms. The last release that ASDM is supported on is 6.1(5).

ASDM Version 6.2 supports the following platforms and releases:

- ASA 5505, software Version 8.2(2), 8.2(1), 8.0(2), 8.0(3), 8.0(4), and 8.0(5)
- ASA 5510, software Version 8.2(2), 8.2(1), 8.0(2), 8.0(3), 8.0(4), and 8.0(5)
- ASA 5520, software Version 8.2(2), 8.2(1), 8.0(2), 8.0(3), 8.0(4), and 8.0(5)
- ASA 5540, software Version 8.2(2), 8.2(1), 8.0(2), 8.0(3), 8.0(4), and 8.0(5)
- ASA 5550, software Version 8.2(2), 8.2(1), 8.0(2), 8.0(3), 8.0(4), and 8.0(5)

- ASA 5580, software Version 8.2(2), 8.2(1), 8.1(1) and 8.1(2)

ASDM Version 6.2 supports the following SSMs and releases:

- Advanced Inspection and Prevention (AIP) SSM, software Versions 5.0, 5.1, 6.0, 6.1, and 6.2
- Content Security and Control (CSC) SSM, software Versions 6.1, 6.2, and 6.3

See the "SSM and SSC Support Per Model" section on page 3-2 for more information.

Multiple ASDM Session Support

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 five concurrent ASDM sessions in single, routed mode. Only one session per browser per PC or workstation is supported for a specified security appliance. In multiple context mode, five concurrent ASDM sessions are supported per context, up to a maximum of 32 total connections for each security appliance.

Unsupported Commands

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

This section includes the following topics:

- Ignored and View-Only Commands, page 2-3
- Effects of Unsupported Commands, page 2-4
- Discontinuous Subnet Masks Not Supported, page 2-5
- Interactive User Commands Not Supported by the ASDM CLI Tool, page 2-5

Ignored and View-Only Commands

Table 2-2 lists commands that ASDM supports in the configuration when added through 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 the command is view-only, then it appears in the GUI, but you cannot edit it.

Table 2-2 List of Unsupported Commands

Unsupported Commands	ASDM Behavior
access-list	Ignored if not used.
capture	Ignored.
coredump	Ignored. This can be configured only using the CLI.
eject	Unsupported.
established	Ignored.
failover timeout	Ignored.
ipv6 nd prefix	Unsupported.

Table 2-2 List of Unsupported Commands **(continued)**

Unsupported Commands	ASDM Behavior
match-metric	Ignored. This is a subcommand of route-map.
match-interface	Ignored. This is a subcommand of route-map.
match route-type	Ignored. This is a subcommand of route-map.
pager	Ignored.
pim accept-register route-map	Ignored. You can configure only the **list** option using ASDM.
prefix-list	Ignored if not used in an OSPF area.

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`

set metric	Ignored.
sysopt nodnsalias	Ignored.
sysopt uauth allow-http-cache	Ignored.
terminal	Ignored.
tunnel-group name general-attributes dhcp-server	The **dhcp-server** subcommand is unsupported. ASDM only allows one setting for all DHCP servers.

Effects of Unsupported Commands

- If ASDM loads an existing running configuration and finds other unsupported commands, ASDM operation is unaffected. To view the unsupported commands, choose Tools > 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 ASA 5500 Series 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 three by your system administrator, which allows Monitor-only mode. For more information, choose Configuration > Device Management > Users/AAA > User Accounts and Configuration > Device Management > Users/AAA > AAA Access.

Discontinuous Subnet Masks Not Supported

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
```

Interactive User Commands Not Supported by the ASDM CLI Tool

The ASDM CLI tool does not support interactive user commands. If you enter a CLI command that requires interactive confirmation, ASDM prompts you to enter “[yes/no]” but does not recognize your input. ASDM then times out waiting for your response.

For example:

1. From the ASDM Tools menu, click **Command Line Interface**.
2. Enter the **crypto key generate rsa** command.

   ASDM generates the default 1024-bit RSA key.
3. Enter the **crypto key generate rsa** command again.

   Instead of regenerating the RSA keys by overwriting the previous one, ASDM displays the following error:

   ```
   Do you really want to replace them? [yes/no]:WARNING: You already have RSA
   ke0000000000000$A key
   Input line must be less than 16 characters in length.

   %Please answer 'yes' or 'no'.
   Do you really want to replace them [yes/no]:

   %ERROR: Timed out waiting for a response.
   ERROR: Failed to create new RSA keys names <Default-RSA-key>
   ```

Workaround:

- You can configure most commands that require user interaction by means of the ASDM panes.
- For CLI commands that have a **noconfirm** option, use this option when entering the CLI command. For example:

  ```
  crypto key generate rsa noconfirm
  ```

Defining Preferences

This feature lets you change the behavior of some ASDM functions between sessions.

To change various settings in ASDM, perform the following steps:

Step 1 In the main ASDM application window, choose **Tools > Preferences**.

The Preferences dialog box appears, with three tabs: General, Rules Table, and Syslog.

Step 2 To define your settings, click one of these tabs: the **General** tab to specify general preferences; the **Rules Tables** tab to specify preferences for the Rules table; and the **Syslog** tab to specify the appearance of syslog messages displayed in the Home pane and to enable the display of a warning message for NetFlow-related syslog messages.

Step 3 On the General tab, specify the following:

a. Check the **Warn that configuration in ASDM is out of sync with the configuration in ASA** check box to be notified when the startup configuration and the running configuration are no longer in sync with each other.

b. Check the **Show configuration restriction message to read-only user** check box to display the following message to a read-only user at startup. This option is checked by default.

```
"You are not allowed to modify the ASA configuration, because you do not have
sufficient privileges."
```

c. Check the **Confirm before exiting ASDM** check box to display a prompt when you try to close ASDM to confirm that you want to exit. This option is checked by default.

d. Check the **Enable screen reader support (requires ASDM restart)** check box to enable screen readers to work. You must restart ASDM to enable this option.

e. Check the **Preview commands before sending them to the device** check box to view CLI commands generated by ASDM.

f. Check the **Enable cumulative (batch) CLI delivery** check box to send multiple commands in a single group to the security appliance.

g. Enter the minimum amount of time in seconds for a configuration to send a timeout message. The default is 60 seconds.

h. To allow the Packet Capture Wizard to display captured packets, enter the name of the network sniffer application or click **Browse** to find it in the file system.

Step 4 On the Rules Tables tab, specify the following:

a. Display settings let you change the way rules appear in the Rules table.

- Check the **Auto-expand network and service object groups with specified prefix** check box to display the network and service object groups automatically expanded based on the Auto-Expand Prefix setting.
- In the Auto-Expand Prefix field, enter the prefix of the network and service object groups to expand automatically when displayed.
- Check the **Show members of network and service object groups** check box to display members of network and service object groups and the group name in the Rules table. If the check box is not checked, only the group name is displayed.
- In the Limit Members To field, enter the number of network and service object groups to display. When the object group members are displayed, then only the first *n* members are displayed.
- Check the **Show all actions for service policy rules** check box to display all actions in the Rules table. When unchecked, a summary appears.

b. Deployment settings let you configure the behavior of the security appliance when deploying changes to the Rules table.

   – Check the **Issue "clear xlate" command when deploying access lists** check box to clear the NAT table when deploying new access lists. This setting ensures the access lists that are configured on the security appliance are applied to all translated addresses.

c. Access Rule Hit Count Settings let you configure the frequency for which the hit counts are updated in the Access Rules table. Hit counts are applicable for explicit rules only. No hit count will be displayed for implicit rules in the Access Rules table.

   – Check the **Update access rule hit counts automatically** check box to have the hit counts automatically updated in the Access Rules table.

   – In the Update Frequency field, specify the frequency in seconds in which the hit count column is updated in the Access Rules table. Valid values are 10 - 86400 seconds.

Step 5 On the Syslog tab, specify the following:

- In the Syslog Colors area, you can customize the message display by configuring background or foreground colors for messages at each severity level. The Severity column lists each severity level by name and number. To change the background color or foreground color for messages at a specified severity level, click the corresponding column. The Pick a Color dialog box appears. Click one of the following tabs:
  – On the Swatches tab, choose a color from the palette, and click **OK**.
  – On the HSB tab, specify the H, S, and B settings, and click **OK**.
  – On the RGB tab, specify the Red, Green, and Blue settings, and click **OK**.
- In the NetFlow area, to enable the display of a warning message to disable redundant syslog messages, check the **Warn to disable redundant syslog messages when NetFlow action is first applied to the global service policy rule** check box.

Step 6 After you have specified settings on these three tabs, click **OK** to save your settings and close the Preferences dialog box.

Note Each time that you check or uncheck a preferences setting, the change is saved to the .conf file and becomes available to all the other ASDM sessions running on the workstation at the time. You must restart ASDM for all changes to take effect.

CHAPTER 3

Introduction to the Security Appliance

The security appliance combines advanced stateful firewall and VPN concentrator functionality in one device, and for some models, an integrated intrusion prevention module called the AIP SSM or an integrated content security and control module called the CSC SSM. The security appliance includes many advanced features, such as multiple security contexts (similar to virtualized firewalls), transparent (Layer 2) firewall or routed (Layer 3) firewall operation, advanced inspection engines, IPSec and clientless SSL support, and many more features.

Note ASDM 6.2(1) and higher is not supported on the PIX platforms. The last release that ASDM is supported on is 6.1(5).

This chapter includes the following sections:

SSM and SSC Support Per Model

Table 3-1 shows the Security Services Modules (SSMs) and Security Services Cards (SSCs) supported by each platform:

Table 3-1 SSM Support

Platform	SSM Models	SSC Models
ASA 5505	No support	AIP SSC 5
ASA 5510	AIP SSM 10 AIP SSM 20 CSC SSM 10 CSC SSM 20 4GE SSM	No support
ASA 5520	AIP SSM 10 AIP SSM 20 AIP SSM 40 CSC SSM 10 CSC SSM 20 4GE SSM	No support
ASA 5540	AIP SSM 10 AIP SSM 20 AIP SSM 40 CSC SSM 10[1] CSC SSM 20[1] 4GE SSM	No support
ASA 5550	No support (the 4GE SSM is built-in and not user-removable)	No support
ASA 5580	No support	No support

[1] The CSC SSM licenses support up to 1000 users while the Cisco ASA 5540 Series appliance can support significantly more users. If you deploy CSC SSM with an ASA 5540 adaptive security appliance, be sure to configure the security appliance to send the CSC SSM only the traffic that should be scanned.

VPN Specifications

See the *Cisco ASA 5500 Series VPN Compatibility Reference* at http://www.cisco.com/en/US/docs/security/asa/compatibility/vpn-platforms-82.html.

New Features

This section includes the following topics:

- New Features in ASDM 6.2(5)/ASA 8.2(2), page 3-3
- New Features in ASDM 6.2(3)/ASA 8.0(5), page 3-5
- New Features in ASDM 6.2(1)/ASA 8.2(1), page 3-6

New Features in ASDM 6.2(5)/ASA 8.2(2)

Table 3-2 lists the new features for ASDM Version 6.2(5). All features apply only to ASA Version 8.2(2).

Table 3-2 New Features for ASDM Version 6.2(5)/ASA Version 8.2(2)

Feature	Description
Remote Access Features	
Scalable Solutions for Waiting-to-Resume VPN Sessions	An administrator can now keep track of the number of users in the active state and can look at the statistics. The sessions that have been inactive for the longest time are marked as idle (and are automatically logged off) so that license capacity is not reached and new users can log in. The following screen was modified: Monitoring > VPN > VPN Statistics > Sessions. *Also available in Version 8.0(5).*
Application Inspection Features	
Inspection for IP Options	You can now control which IP packets with specific IP options should be allowed through the security appliance. You can also clear IP options from an IP packet, and then allow it through the security appliance. Previously, all IP options were denied by default, except for some special cases. **Note** This inspection is enabled by default. Therefore, the security appliance allows RSVP traffic that contains packets with the Router Alert option (option 20) when the security appliance is in routed mode. The following screens were introduced: Configuration > Firewall > Objects > Inspect Maps > IP-Options Configuration > Firewall > Service Policy > Add/Edit Service Policy Rule > Rule Actions > Protocol Inspection
Enabling Call Set up Between H.323 Endpoints	You can enable call setup between H.323 endpoints when the Gatekeeper is inside the network. The security appliance includes options to open pinholes for calls based on the RegistrationRequest/RegistrationConfirm (RRQ/RCF) messages. Because these RRQ/RCF messages are sent to and from the Gatekeeper, the calling endpoint IP address is unknown and the security appliance opens a pinhole through source IP address/port 0/0. By default, this option is disabled. The following screen was modified: Configuration > Firewall > Objects > Inspect Maps > H.323 > Details > State Checking. *Also available in Version 8.0(5).*
Unified Communication Features	

Table 3-2 New Features for ASDM Version 6.2(5)/ASA Version 8.2(2) (continued)

Feature	Description
Mobility Proxy application no longer requires Unified Communications Proxy license	The Mobility Proxy no longer requires the UC Proxy license.
Interface Features	
In multiple context mode, auto-generated MAC addresses now use a user-configurable prefix, and other enhancements	The MAC address format was changed to allow use of a prefix, to use a fixed starting value (A2), and to use a different scheme for the primary and secondary unit MAC addresses in a failover pair. The MAC addresesss are also now persistent accross reloads. The command parser now checks if auto-generation is enabled; if you want to also manually assign a MAC address, you cannot start the manual MAC address with A2. The following screen was modified: Configuration > Context Management > Security Contexts. *Also available in Version 8.0(5).*
Support for Pause Frames for Flow Control on the ASA 5580 10 Gigabit Ethernet Interfaces	You can now enable pause (XOFF) frames for flow control. The following screens were modified: (Single Mode) Configuration > Device Setup > Interfaces > Add/Edit Interface > General (Multiple Mode, System) Configuration > Interfaces > Add/Edit Interface
Firewall Features	
Botnet Traffic Filter Enhancements	The Botnet Traffic Filter now supports automatic blocking of blacklisted traffic based on the threat level. You can also view the category and threat level of malware sites in statistics and reports. Reporting was enhanced to show infected hosts. The 1 hour timeout for reports for top hosts was removed; there is now no timeout. The following screens were introduced or modified: Configuration > Firewall > Botnet Traffic Filter > Traffic Settings Monitoring > Botnet Traffic Filter > Infected Hosts
Routing Features	
DHCP RFC compatibility (rfc3011, rfc3527) to resolve routing issues	This enhancement introduces security appliance support for DHCP RFCs 3011 (The IPv4 Subnet Selection Option) and 3527 (Link Selection Sub-option for the Relay Agent Information Option). For each DHCP server configured for VPN clients, you can now configure the security appliance to send the Subnet Selection option or the Link Selection option. The following screen was modified: Remote Access VPN > Network Access > IPsec connection profiles > Add/Edit. *Also available in Version 8.0(5).*
High Availablility Features	
IPv6 Support in Failover Configurations	IPv6 is now supported in failover configurations. You can assign active and standby IPv6 addresses to interfaces and use IPv6 addresses for the failover and Stateful Failover interfaces. The following screens were modified: Configuration > Device Management > High Availability > Failover > Setup Configuration > Device Management > High Availability > Failover > Interfaces Configuration > Device Management > High Availability > HA/Scalability Wizard

Table 3-2 New Features for ASDM Version 6.2(5)/ASA Version 8.2(2) (continued)

Feature	Description
No notifications when interfaces are brought up or brought down during a switchover event	To distinguish between link up/down transitions during normal operation from link up/down transitions during failover, no link up/link down traps are sent during a failover. Also, no syslog messages about link up/down transitions during failover are sent. *Also available in Version 8.0(5).*
AAA Features	
100 AAA Server Groups	You can now configure up to 100 AAA server groups; the previous limit was 15 server groups. The following screen was modified: Configuration > Device Management > Users/AAA > AAA Server Groups.
Monitoring Features	
Smart Call Home	Smart Call Home offers proactive diagnostics and real-time alerts on the security appliance and provides higher network availability and increased operational efficiency. Customers and TAC engineers get what they need to resolve problems quickly when an issue is detected. **Note** Smart Call Home server Version 3.0(1) has limited support for the security appliance. See the “Important Notes” for more information. The following screen was introduced: Configuration> Device Management> Smart Call Home.

New Features in ASDM 6.2(3)/ASA 8.0(5)

Table 3-3 lists the new features for ASDM Version 6.2(3). All features apply only to ASA Version 8.0(5).

Note Version 8.0(5) is not supported on the PIX security appliance.

Table 3-3 New Features for ASDM Version 6.2(3)/ASA Version 8.0(5)

Feature	Description
Remote Access Features	
Scalable Solutions for Waiting-to-Resume VPN Sessions	An administrator can now keep track of the number of users in the active state and can look at the statistics. The sessions that have been inactive for the longest time are marked as idle (and are automatically logged off) so that license capacity is not reached and new users can log in The following ASDM screen was modified: Monitoring > VPN > VPN Statistics > Sessions.
Application Inspection Features	
Enabling Call Set up Between H.323 Endpoints	You can enable call setup between H.323 endpoints when the Gatekeeper is inside the network. The security appliance includes options to open pinholes for calls based on the RegistrationRequest/RegistrationConfirm (RRQ/RCF) messages. Because these RRQ/RCF messages are sent to and from the Gatekeeper, the calling endpoint's IP address is unknown and the security appliance opens a pinhole through source IP address/port 0/0. By default, this option is disabled. The following ASDM screen was modified: Configuration > Firewall > Objects > Inspect Maps > H.323 > Details > State Checking.

Table 3-3 New Features for ASDM Version 6.2(3)/ASA Version 8.0(5) (continued)

Feature	Description
Interface Features	
In multiple context mode, auto-generated MAC addresses now use a user-configurable prefix, and other enhancements	The MAC address format was changed to allow use of a prefix, to use a fixed starting value (A2), and to use a different scheme for the primary and secondary unit MAC addresses in a failover pair. The MAC addresess are also now persistent accross reloads. The command parser now checks if auto-generation is enabled; if you want to also manually assign a MAC address, you cannot start the manual MAC address with A2. The following ASDM screen was modified: Configuration > Context Management > Security Contexts.
High Availablility Features	
No notifications when interfaces are brought up or brought down during a switchover event	To distinguish between link up/down transitions during normal operation from link up/down transitions during failover, no link up/link down traps are sent during a failover. Also, no syslog messages about link up/down transitions during failover are sent.
Routing Features	
DHCP RFC compatibility (rfc3011, rfc3527) to resolve routing issues	This enhancement introduces security appliance support for DHCP RFCs 3011 (The IPv4 Subnet Selection Option) and 3527 (Link Selection Sub-option for the Relay Agent Information Option). The following ASDM screen was modified: Remote Access VPN > Network Access > IPsec connection profiles > Add/Edit.
SSM Features	
CSC 6.3 Support in ASDM	ASDM displays Web Reputation, User Group Policies, and User ID Settings in the Plus License listing on the main home page. CSC 6.3 security event enhancements are included, such as the new Web Reputation events and user and group identifications.

New Features in ASDM 6.2(1)/ASA 8.2(1)

Table 3-4 lists the new features for ASDM Version 6.2(1). All features apply only to Version 8.2(1).

Table 3-4 New Features for ASDM Version 6.2(1)/ASA Version 8.2(1)

Feature	Description
Remote Access Features	
One Time Password Support for ASDM Authentication	ASDM now supports administrator authentication using one time passwords (OTPs) supported by RSA SecurID (SDI). This feature addresses security concerns about administrators authenticating with static passwords. New session controls for ASDM users include the ability to limit the session time and the idle time. When the password used by the ASDM administrator times out, ASDM prompts the administrator to re-authenticate. In ASDM, see Configuration > Device Management > Management Access > ASDM/HTTPD/Telnet/SSH.

Table 3-4 New Features for ASDM Version 6.2(1)/ASA Version 8.2(1) (continued)

Feature	Description
Customizing Secure Desktop	You can use ASDM to customize the Secure Desktop windows displayed to remote users, including the Secure Desktop background (the lock icon) and its text color, and the dialog banners for the Desktop, Cache Cleaner, Keystroke Logger, and Close Secure Desktop windows. In ASDM, see Configuration > CSD Manager > Secure Desktop Manager.
Pre-fill Username from Certificate	The pre-fill username feature enables the use of a username extracted from a certificate for username/password authentication. With this feature enabled, the username is "pre-filled" on the login screen, with the user being prompted only for the password. The double-authentication feature is compatible with the pre-fill username feature, as the pre-fill username feature can support extracting a primary username and a secondary username from the certificate to serve as the usernames for double authentication when two usernames are required. In ASDM, see Configuration> Remote Access VPN > Network (Client) Access > AnyConnect or Clienltess SSL VPN Connection Profiles > Advanced. Settings are in the Authentication, Secondary Authentication, and Authorization panes.
Double Authentication	The double authentication feature implements two-factor authentication for remote access to the network, in accordance with the Payment Card Industry Standards Council Data Security Standard. This feature requires that the user enter two separate sets of login credentials at the login page. For example, the primary authentication might be a one-time password, and the secondary authentication might be a domain (Active Directory) credential. If either authentication fails, the connection is denied. Both the AnyConnect VPN client and Clientless SSL VPN support double authentication. The AnyConnect client supports double authentication on Windows computers (including supported Windows Mobile devices and Start Before Logon), Mac computers, and Linux computers. The IPsec VPN client, SVC client, cut-through-proxy authentication, hardware client authentication, and management authentication do not support double authentication. **Note** The RSA/SDI authentication server type cannot be used as the secondary username/password credential. It can only be used for primary authentication. In ASDM, see Configuration > Remote Access VPN > Network (Client) Access or Clientless SSL VPN > AnyConnect Connection Profiles > Add/Edit > Advanced > Secondary Authentication.
AnyConnect Essentials	AnyConnect Essentials is a separately licensed SSL VPN client, entirely configured on the security appliance, that provides the full AnyConnect capability, with the following exceptions: • No CSD (including HostScan/Vault/Cache Cleaner) • No clientless SSL VPN • Optional Windows Mobile Support The AnyConnect Essentials client provides remote end users running Microsoft Windows Vista, Windows Mobile, Windows XP or Windows 2000, Linux, or Macintosh OS X, with the benefits of a Cisco SSL VPN client. **Note** This license cannot be used at the same time as the shared SSL VPN premium license. In ASDM, see Configuration > Remote Access VPN > Network (Client) Access > Advanced > AnyConnect Essentials License. The AnyConnect Essentials license must be installed for ASDM to show this pane.

Table 3-4 New Features for ASDM Version 6.2(1)/ASA Version 8.2(1) (continued)

Feature	Description
Disabling Cisco Secure Desktop per Connection Profile	When enabled, Cisco Secure Desktop automatically runs on all computers that make SSL VPN connections to the security appliance. This new feature lets you exempt certain users from running Cisco Secure Desktop on a per connection profile basis. It prevents the detection of endpoint attributes for these sessions, so you might need to adjust the Dynamic Access Policy (DAP) configuration. In ASDM, see Configuration > Remote Access VPN > Clientless SSL VPN Access > Connection Profiles > Add or Edit > Advanced, Clientless SSL VPN Configuration. or Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Connection Profiles > Add or Edit > Advanced > SSL VPN.
Certificate Authentication Per Connection Profile	Previous versions supported certificate authentication for each security appliance interface, so users received certificate prompts even if they did not need a certificate. With this new feature, users receive a certificate prompt only if the connection profile configuration requires a certificate. This feature is automatic. In ASDM, see Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Connection Profiles > Add/Edit > Basic. or Configuraiton > Remote Access VPN > Clientless SSL VPN > Connection Profiles > Add/Edit>Basic.
EKU Extensions for Certificate Mapping	This feature adds the ability to create certificate maps that look at the Extended Key Usage extension of a client certificate and use these values in determining what connection profile the client should use. If the client does not match that profile, it uses the default group. The outcome of the connection then depends on whether or not the certificate is valid and the authentication settings of the connection profile. In ASDM, use the IPSec Certificate to Connection Maps > Rules pane, or Certificate to SSL VPN Connections Profile Maps pane.
SSL VPN SharePoint Support for Win 2007 Server	Clientless SSL VPN sessions now support Microsoft Office SharePoint Server 2007.
Shared license for SSL VPN sessions	You can purchase a shared license with a large number of SSL VPN sessions and share the sessions as needed among a group of security appliances by configuring one of the security appliances as a shared license server, and the rest as clients. **Note** This license cannot be used at the same time as the AnyConnect Essentials license. In ASDM, see Configuration > Device Management > Licensing > Shared SSL VPN Licenses. Also see, Monitoring > VPN > Clientless SSL VPN > Shared Licenses.
Updated VPN Wizard	The VPN Wizard (accessible by choosing Wizards > IPSec VPN Wizard) was updated. The step to select IPsec Encryption and Authentication (formerly Step 9 of 11) was removed because the Wizard now generates default values for these settings. In addition, the step to select IPsec Settings (Optional) now includes new fields to enable perfect forwarding secrecy (PFS) and set the Diffie-Hellman Group.
Firewall Features	

Table 3-4 New Features for ASDM Version 6.2(1)/ASA Version 8.2(1) (continued)

Feature	Description
TCP state bypass	If you have asymmetric routing configured on upstream routers, and traffic alternates between two security appliances, then you can configure TCP state bypass for specific traffic. In ASDM, see Configuration > Firewall > Service Policy Rules > Rule Actions > Connection Settings.
Per-Interface IP Addresses for the Media-Termination Instance Used by the Phone Proxy	In Version 8.0(4), you configured a global media-termination address (MTA) on the security appliance. In Version 8.2, you can now configure MTAs for individual interfaces (with a minimum of two MTAs). As a result of this enhancement, the old CLI has been deprecated. You can continue to use the old configuration if desired. However, if you need to change the configuration at all, only the new configuration method is accepted; you cannot later restore the old configuration. In ASDM, see Configuration > Firewall > Advanced > Encrypted Traffic Inspection > Media Termination Address.
H.239 Message Support in H.323 Application Inspection	In this release, the security appliance supports the H.239 standard as part of H.323 application inspection. H.239 is a standard that provides the ability for H.300 series endpoints to open an additional video channel in a single call. In a call, an endpoint (such as a video phone), sends a channel for video and a channel for data presentation. The H.239 negotiation occurs on the H.245 channel. The security appliance opens a pinhole for the additional media channel. The endpoints use open logical channel message (OLC) to signal a new channel creation. The message extension is part of H.245 version 13. The decoding and encoding of the telepresentation session is enabled by default. H.239 encoding and decoding is preformed by ASN.1 coder. In ASDM, see Configuration > Firewall > Service Policy Rules > Add Service Policy Rule Wizard > Rule Actions > Protocol Inspection > H.323 H.225. Click **Configure** and then choose the H.323 Inspect Map.
Processing H.323 Endpoints When the Endpoints Do Not Send OLCAck	H.323 application inspection has been enhanced to process common H.323 endpoints. The enhancement affects endpoints using the extendedVideoCapability OLC with the H.239 protocol identifier. Even when an H.323 endpoint does not send OLCAck after receiving an OLC message from a peer, the security appliance propagates OLC media proposal information into the media array and opens a pinhole for the media channel (extendedVideoCapability). In ASDM, see Configuration > Firewall > Service Policy Rules > Add Service Policy Rule Wizard > Rule Actions > Protocol Inspection > H.323 H.225.
IPv6 in transparent firewall mode	Transparent firewall mode now participates in IPv6 routing. Prior to this release, the security appliance could not pass IPv6 traffic in transparent mode. You can now configure an IPv6 management address in transparent mode, create IPv6 access lists, and configure other IPv6 features; the security appliance recognizes and passes IPv6 packets. All IPv6 functionality is supported unless specifically noted. In ASDM, see Configuration > Device Management > Management Access > Management IP Address.

Table 3-4 New Features for ASDM Version 6.2(1)/ASA Version 8.2(1) (continued)

Feature	Description
Botnet Traffic Filter	Malware is malicious software that is installed on an unknowing host. Malware that attempts network activity such as sending private data (passwords, credit card numbers, key strokes, or proprietary data) can be detected by the Botnet Traffic Filter when the malware starts a connection to a known bad IP address. The Botnet Traffic Filter checks incoming and outgoing connections against a dynamic database of known bad domain names and IP addresses, and then logs any suspicious activity. You can also supplement the dynamic database with a static database by entering IP addresses or domain names in a local "blacklist" or "whitelist." **Note** This feature requires the Botnet Traffic Filter license. See the following licensing document for more information: http://www.cisco.com/en/US/docs/security/asa/asa82/license/license82.html In ASDM, see Configuration > Firewall > Botnet Traffic Filter.
AIP SSC card for the ASA 5505	The AIP SSC offers IPS for the ASA 5505 security appliance. Note that the AIP SSM does not support virtual sensors. In ASDM, see Configuration > Device Setup > SSC Setup and Configuration > IPS.
IPv6 support for IPS	You can now send IPv6 traffic to the AIP SSM or SSC when your traffic class uses the **match any** command, and the policy map specifies the **ips** command. In ASDM, see Configuration > Firewall > Service Policy Rules.
Management Features	
SNMP version 3 and encryption	This release provides DES, 3DES, or AES encryption and support for SNMP Version 3, the most secure form of the supported security models. This version allows you to configure authentication characteristics by using the User-based Security Model (USM). In ASDM, see Configuration > Device Management > Management Access > SNMP.
Routing Features	
Multicast NAT	The security appliance now offers Multicast NAT support for group addresses.
Troubleshooting Features	
Coredump functionality	A coredump is a snapshot of the running program when the program has terminated abnormally. Coredumps are used to diagnose or debug errors and save a crash for later or off-site analysis. Cisco TAC may request that users enable the coredump feature to troubleshoot application or system crashes on the security appliance. To enable coredump, use the **coredump enable** command.
ASDM Features	
ASDM Support for IPv6	All IPv6 functionality is supported unless specifically noted.
Support for Public Server configuration	You can use ASDM to configure a public server. This allows to you define servers and services that you want to expose to an outside interface. In ASDM, see Configuration > Firewall > Public Servers.

Firewall Functional Overview

Firewalls protect inside networks from unauthorized access by users on an outside network. A firewall can also protect inside networks from each other, for example, by keeping a human resources network separate from a user network. If you have network resources that need to be available to an outside user, such as a web or FTP server, you can place these resources on a separate network behind the firewall, called a *demilitarized zone* (DMZ). The firewall allows limited access to the DMZ, but because the DMZ only includes the public servers, an attack there only affects the servers and does not affect the other inside networks. You can also control when inside users access outside networks (for example, access to the Internet), by allowing only certain addresses out, by requiring authentication or authorization, or by coordinating with an external URL filtering server.

When discussing networks connected to a firewall, the *outside* network is in front of the firewall, the *inside* network is protected and behind the firewall, and a *DMZ*, while behind the firewall, allows limited access to outside users. Because the security appliance lets you configure many interfaces with varied security policies, including many inside interfaces, many DMZs, and even many outside interfaces if desired, these terms are used in a general sense only.

This section includes the following topics:

- Configuring Public Servers, page 3-11
- Security Policy Overview, page 3-15
- Firewall Mode Overview, page 3-17
- Stateful Inspection Overview, page 3-17

Configuring Public Servers

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

- Public Server Overview, page 3-11
- Add a Public Server, page 3-12
- Edit a Public Server, page 3-13

Public Server Overview

While one of the basic functions of a firewall is to protect inside networks from unauthorized access by users on an outside network, or protect inside networks from each other, this function involves multiple configurations. That is, inside the DMZ interfaces, creating ACL lists and rules, NAT/PAT rules, and application inspection.

ASDM provides the Public Servers pane in the Configuration > Firewall > Public Servers pane, so that an adminstrator can publish various application servers to be accessed by internal and external users. When selected, this pane displays a list of public servers. internal and external addresses, the interfaces that the internal or external addresses apply to, and the service that is exposed.

In this pane you can add, edit, delete, or modify existing public servers.

Fields

- Add—Adds a public server.

- Edit—Edits a a public server group.
- Delete—Deletes a specified public server.
- Apply—Applies the changes that have been made.
- Reset—Resets the security appliance to the previous 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 a Public Server

To add a public server, perform the following steps:

Note STATIC PAT is not supported on a public server.

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

The Add Public Server dialog box appears.

Step 2 Fill in the values of the Private Interface, Private IP Address, Service, Public Interface, and Public IP Address.

- Private Interface—Use the drop down menu to select the name of the private interface or enter the name in the field.
- Private IP Address—Click the **...** browse button next to the Private IP address field to select the private IP address. The Browse Private IP Address dialog box appears.
- You can enter a name or Private IP address in the Filter field and click Filter. The wildcard characters asterisk (*) and question mark (?) are allowed. To delete the name you just typed, click Clear, or you can click Add. The Add Network Object dialog box appears.

Step 3 Fill in the following values in the Add Network Object dialog box:

- Name—(Optional) The object name. Use characters a to z, A to Z, 0 to 9, a dot, a dash, or an underscore. The name must be 64 characters or less.
- IP Address—The IP address (host address).
- Netmask—The subnet mask for the IP address.
- Description—(Optional) The description of the network object.

Step 4 Click OK.

You can now use this network object when you create a rule. For an edited object, the change is inherited automatically by any rules using the object.

Note You cannot delete a network object that is in use.

Step 5 Service—The service that is exposed to the outside. You can choose any of the currently defined servies or a servie group that has been created. Multiple services from various ports can be openedto the outside.

Click the **...** browse button next to the Service address field to select the service. The Browse Service dialog box appears.

Browse Service Groups dialog box lets you choose a service group. This dialog box is used in multiple configuration windows, and is named appropriately for your current task.

Step 6 In the Browse Service Groups dialog box, choose the Service from the main menu, then click OK. Fill in the following values:

- Public Interface—Use the drop down menu to choose the name of the public interface or enter the name in the field.
- Public IP Address—The address of the server as seen from the outside. If IPv6 has been enabled, the list of IPv6 addresses will be visible from this list.
  Click the **...** browse button next to the Public IP address field to choose the public IP address. The Browse Public IP Address dialog box appears.

Step 7 You can enter a name or Public IP address in the Filter field and click Filter. The wildcard characters asterisk (*) and question mark (?) are allowed. To delete the name you just typed, click Clear, or you can click Add. The Add Network Object dialog box appears.

Step 8 Fill in the following values from the Add Network Object dialog:

- Name—(Optional) The object name. Use characters a to z, A to Z, 0 to 9, a dot, a dash, or an underscore. The name must be 64 characters or less.
- IP Address—The IP address (host address).
- Netmask—The subnet mask for the IP address.
- Description—(Optional) The description of the network object.

Step 9 Click OK.

You can now use this network object when you create a rule. For an edited object, the change is inherited automatically by any rules using the object.

Note Public Server rules are not be applicable for host address that are used as network-object group member in an access list.
For example:

```
# object-group network k1
# network-object 10.16.1.1 255.255.255.255
# access-list outside_access_in permit tcp any object-group k1
# access-group outside_access_in in interface outside
# static (inside,outside) 192.168.1.1 10.16.1.1 netmask 255.255.255.255
```

Edit a Public Server

To configure a public server, perform the following steps:

Step 1 In the Configuration > Firewall > Public Servers pane, click **Edit** to edit an object, or choose an existing public server and and click **Edit**.

The Edit Public Server dialog box appears.

Step 2 Fill in the values of the Inside Interface, Inside Address, Service, Outside Interface, Outside Address by performing the following steps.

- Private Interface—Lists the inside interfaces that are currently defined and provides the interface where the server is located.
  Use the pulldown menu to select the name of the interface or enter the name in the field.
- Private IP Address—The address of the server as seen from the outside. If IPv6 has been enabled, the list of IPv6 addresses will be visible from this list
  Click the **...** browse button next to the Private IP address field to select the private IP address. The Browse Private IP Address dialog box appears.
- You can enter a name or inside IP address in the Filter field and click Filter. The wildcard characters asterisk (*) and question mark (?) are allowed. To delete the name you just typed, click **Clear**, or you can click Add. The Add Network Object dialog box appears.

Step 3 Fill in the following values from the Add Network Object dialog:

- Name—(Optional) The object name. Use characters a to z, A to Z, 0 to 9, a dot, a dash, or an underscore. The name must be 64 characters or less.
- IP Address—The IP address (host address).
- Netmask—The subnet mask for the IP address.
- Description—(Optional) The description of the network object.

Step 4 Click OK.

You can now use this network object when you create a rule. For an edited object, the change is inherited automatically by any rules using the object.

Note You cannot delete a network object that is in use.

Step 5 Service—The service that is exposed to the outside. You can choose any of the currently defined servies or a servie group that has been created. Multiple services from various ports can be openedto the outside.

Click the **...** browse button next to the Service address field to choose the service.
The Browse Service dialog box appears.

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.

Step 6 On the Browse Service Groups dialog box, choose the Service from the main menu, then click OK.

- Public Interface—A drop down list that displays the interfaces that are currently defined and allows tou to specify which interface has access to the server.

  Use the pulldown menu to select the name of the outside interface or enter the name in the field.
- Public IP Address—Click the **...** browse button next to the Public IP Address field to select the outside IP address. The Browse IP Address dialog box appears.
- You can enter a name or IP address in the Filter field and click Filter. The wildcard characters asterisk (*) and question mark (?) are allowed. To delete the name you just typed, click **Clear**, or you can click Add. The Add Network Object dialog box appears.

Step 7 Fill in the following values from the Add Network Object dialog:

- Name—(Optional) The object name. Use characters a to z, A to Z, 0 to 9, a dot, a dash, or an underscore. The name must be 64 characters or less.

- IP Address—The IP address (host address).
- Netmask—The subnet mask for the IP address.
- Description—(Optional) The description of the network object.

Step 8 Click OK.

You can now use this network object when you create a rule. For an edited object, the change is inherited automatically by any rules using the object.

Security Policy Overview

A security policy determines which traffic is allowed to pass through the firewall to access another network. By default, the security appliance allows traffic to flow freely from an inside network (higher security level) to an outside network (lower security level). You can apply actions to traffic to customize the security policy. This section includes the following topics:

- Permitting or Denying Traffic with Access Lists, page 3-15
- Applying NAT, page 3-15
- Protecting from IP Fragments, page 3-16
- Using AAA for Through Traffic, page 3-16
- Applying HTTP, HTTPS, or FTP Filtering, page 3-16
- Applying Application Inspection, page 3-16
- Sending Traffic to the Advanced Inspection and Prevention Security Services Module, page 3-16
- Sending Traffic to the Content Security and Control Security Services Module, page 3-16
- Applying QoS Policies, page 3-16
- Applying Connection Limits and TCP Normalization, page 3-17

Permitting or Denying Traffic with Access Lists

You can apply an access list to limit traffic from inside to outside, or allow traffic from outside to inside. For transparent firewall mode, you can also apply an EtherType access list to allow non-IP traffic.

Applying NAT

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.

Protecting from IP Fragments

The security appliance provides IP fragment protection. This feature performs full reassembly of all ICMP error messages and virtual reassembly of the remaining IP fragments that are routed through the security appliance. Fragments that fail the security check are dropped and logged. Virtual reassembly cannot be disabled.

Using AAA for Through Traffic

You can require 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.

Applying HTTP, HTTPS, or FTP Filtering

Although you can use access lists to prevent outbound access to specific websites or FTP servers, configuring and managing web usage this way is not practical because of the size and dynamic nature of the Internet. We recommend that you use the security appliance in conjunction with a separate server running one of the following Internet filtering products:

- Websense Enterprise
- Secure Computing SmartFilter

Applying Application Inspection

Inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the security appliance to perform a deep packet inspection.

Sending Traffic to the Advanced Inspection and Prevention Security Services Module

If your model supports the AIP SSM for intrusion prevention, then you can send traffic to the AIP SSM for inspection. The AIP SSM is an intrusion prevention services module that monitors and performs real-time analysis of network traffic by looking for anomalies and misuse based on an extensive, embedded signature library. When the system detects unauthorized activity, it can terminate the specific connection, permanently block the attacking host, log the incident, and send an alert to the device manager. Other legitimate connections continue to operate independently without interruption. For more information, see *Configuring the Cisco Intrusion Prevention System Sensor Using the Command Line Interface*.

Sending Traffic to the Content Security and Control Security Services Module

If your model supports it, 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 the security appliance to send to it.

Applying QoS Policies

Some network traffic, such as voice and streaming video, cannot tolerate long latency times. QoS is a network feature that lets you give priority to these types of traffic. QoS refers to the capability of a network to provide better service to selected network traffic.

Applying Connection Limits and TCP Normalization

You can limit TCP and UDP connections and embryonic connections. 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 normalization is a feature consisting of advanced TCP connection settings designed to drop packets that do not appear normal.

Enabling Threat Detection

You can configure scanning threat detection and basic threat detection, and also how to use statistics to analyze threats.

Basic threat detection detects activity that might be related to an attack, such as a DoS attack, and automatically sends a system log message.

A typical scanning attack consists of a host that tests the accessibility of every IP address in a subnet (by scanning through many hosts in the subnet or sweeping through many ports in a host or subnet). The scanning threat detection feature determines when a host is performing a scan. Unlike IPS scan detection that is based on traffic signatures, the security appliance scanning threat detection feature maintains an extensive database that contains host statistics that can be analyzed for scanning activity.

The host database tracks suspicious activity such as connections with no return activity, access of closed service ports, vulnerable TCP behaviors such as non-random IPID, and many more behaviors.

You can configure the security appliance to send system log messages about an attacker or you can automatically shun the host.

Firewall Mode Overview

The security appliance runs in two different firewall modes:

- Routed
- Transparent

In routed mode, the security appliance is considered to be a router hop in the network.

In transparent mode, the security appliance acts like a “bump in the wire,” or a “stealth firewall,” and is not considered a router hop. The security appliance connects to the same network on its inside and outside interfaces.

You might use a transparent firewall to simplify your network configuration. Transparent mode is also useful if you want the firewall to be invisible 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 using an EtherType access list.

Stateful Inspection Overview

All traffic that goes through the security appliance is inspected using the Adaptive Security Algorithm and either allowed through or dropped. A simple packet filter can check for the correct source address, destination address, and ports, but it does not check that the packet sequence or flags are correct. A filter also checks every packet against the filter, which can be a slow process.

A stateful firewall like the security appliance, however, takes into consideration the state of a packet:

- Is this a new connection?

  If it is a new connection, the security appliance has to check the packet against access lists and perform other tasks to determine if the packet is allowed or denied. To perform this check, the first packet of the session goes through the “session management path,” and depending on the type of traffic, it might also pass through the “control plane path.”

  The session management path is responsible for the following tasks:

  – Performing the access list checks

  – Performing route lookups

  – Allocating NAT translations (xlates)

  – Establishing sessions in the “fast path”

Note The session management path and the fast path make up the “accelerated security path.”

  Some packets that require Layer 7 inspection (the packet payload must be inspected or altered) are passed on to the control plane path. Layer 7 inspection engines are required for protocols that have two or more channels: a data channel, which uses well-known port numbers, and a control channel, which uses different port numbers for each session. These protocols include FTP, H.323, and SNMP.

- Is this an established connection?

  If the connection is already established, the security appliance does not need to re-check packets; most matching packets can go through the fast path in both directions. The fast path is responsible for the following tasks:

  – IP checksum verification

  – Session lookup

  – TCP sequence number check

  – NAT translations based on existing sessions

  – Layer 3 and Layer 4 header adjustments

  For UDP or other connectionless protocols, the security appliance creates connection state information so that it can also use the fast path.

  Data packets for protocols that require Layer 7 inspection can also go through the fast path.

  Some established session packets must continue to go through the session management path or the control plane path. Packets that go through the session management path include HTTP packets that require inspection or content filtering. Packets that go through the control plane path include the control packets for protocols that require Layer 7 inspection.

VPN Functional Overview

A VPN is a secure connection across a TCP/IP network (such as the Internet) that appears as a private connection. This secure connection is called a tunnel. 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 invokes various standard protocols to accomplish these functions.

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

The security appliance invokes various standard protocols to accomplish these functions.

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 into the admin context, then that user has system administrator rights and can access the system and all other contexts.

Note You can run all your contexts in routed mode or transparent mode; you cannot run some contexts in one mode and others in another.

Multiple context mode supports static routing only.

CHAPTER 4

Managing Feature Licenses

A license specifies the options that are enabled on a given security appliance. It is represented by an activation key which is a 160-bit (5 32-bit words or 20 bytes) value. This value encodes the serial number (an 11 character string) and the enabled features.

This chapter describes how to obtain an activation key and activate it. It also describes the available licenses for each model. This chapter includes the following sections:

- Supported Feature Licenses Per Model, page 4-1
- Information About Feature Licenses, page 4-9
- Guidelines and Limitations, page 4-17
- Viewing Your Current License, page 4-18
- Obtaining an Activation Key, page 4-18
- Entering a New Activation Key, page 4-19
- Upgrading the License for a Failover Pair, page 4-20
- Configuring a Shared License, page 4-22
- Feature History for Licensing, page 4-24

Supported Feature Licenses Per Model

This section describes the licenses available for each model as well as important notes about licenses. This section includes the following topics:

- Licenses Per Model, page 4-1
- License Notes, page 4-8

Licenses Per Model

This section lists the feature licenses available for each model:

- ASA 5505, Table 4-1 on page 4-2
- ASA 5510, Table 4-2 on page 4-3
- ASA 5520, Table 4-3 on page 4-4
- ASA 5540, Table 4-4 on page 4-5

- ASA 5550, Table 4-5 on page 4-6
- ASA 5580, Table 4-6 on page 4-7

Items that are in italics are separate, optional licenses with which that you can replace the Base or Security Plus license. You can mix and match licenses, for example, the 10 security context license plus the Strong Encryption license; or the 500 Clientless SSL VPN license plus the GTP/GPRS license; or all four licenses together.

Table 4-1 ASA 5505 Adaptive Security Appliance License Features

ASA 5505	Base License			Security Plus		
Firewall Licenses						
Botnet Traffic Filter	Disabled	*Optional temporary license: Available*		Disabled	*Optional temporary license: Available*	
Firewall Conns, Concurrent	10 K			25 K		
GTP/GPRS	No support			No support		
Unified Comm. Sessions[1]	2	*Optional license: 24*		2	*Optional license: 24*	
VPN Licenses						
Adv. Endpoint Assessment	Disabled	*Optional license: Available*		Disabled	*Optional license: Available*	
AnyConnect Essentials[1]	Disabled	*Optional license: Available*		Disabled	*Optional license: Available*	
AnyConnect Mobile[1]	Disabled	*Optional license: Available*		Disabled	*Optional license: Available*	
AnyConnect Premium SSL VPN (sessions)[1]	2	*Optional Permanent licenses:*		2	*Optional Permanent licenses:*	
		10	*25*		*10*	*25*
IPSec VPN (sessions)[1]	10 (max. 25 combined IPSec and SSL VPN)			25 (max. 25 combined IPSec and SSL VPN)		
VPN Load Balancing	No support			No support		
General Licenses						
Encryption	Base (DES)	*Opt. lic.: Strong (3DES/AES)*		Base (DES)	*Opt. lic.: Strong (3DES/AES)*	
Failover	No support			Active/Standby (no stateful failover)		
Security Contexts	No support			No support		
Users, concurrent[2]	10[3]	*Optional licenses:*		10[3]	*Optional licenses:*	
		50	*Unlimited*		*50*	*Unlimited*
VLANs/Zones, Maximum	3 (2 regular zones and 1 restricted zone)			20		
VLAN Trunk, Maximum	No support			8 trunks		

[1] See the "License Notes" section.

[2] In routed mode, hosts on the inside (Business and Home VLANs) count towards the limit only when they communicate with the outside (Internet VLAN). Internet hosts are not counted towards the limit. Hosts that initiate traffic between Business and Home are also not counted towards the limit. The interface associated with the default route is considered to be the Internet interface. If there is no default route, hosts on all interfaces are counted toward the limit. In transparent mode, the interface with the lowest number of hosts is counted towards the host limit. See the **show local-host** command to view host limits.

[3] For a 10-user license, the max. DHCP clients is 32. For 50 users, the max. is 128. For unlimited users, the max. is 250, which is the max. for other models.

Table 4-2 ASA 5510 Adaptive Security Appliance License Features

ASA 5510	Base License						Security Plus					
Firewall Licenses												
Botnet Traffic Filter	Disabled		*Optional temporary license: Available*				Disabled		*Optional temporary license: Available*			
Firewall Conns, Concurrent	50 K						130 K					
GTP/GPRS	No support						No support					
Unified Comm. Sessions[1]	2	*Optional licenses:*					2	*Optional licenses:*				
		24	*50*	*100*				*24*	*50*	*100*		
VPN Licenses												
Adv. Endpoint Assessment	Disabled		*Optional license: Available*				Disabled		*Optional license: Available*			
AnyConnect Essentials[1]	Disabled		*Optional license: Available*				Disabled		*Optional license: Available*			
AnyConnect Mobile[1]	Disabled		*Optional license: Available*				Disabled		*Optional license: Available*			
AnyConnect Premium SSL VPN (sessions)[1]	2	*Optional Permanent licenses:*					2	*Optional Permanent licenses:*				
		10	*25*	*50*	*100*	*250*		*10*	*25*	*50*	*100*	*250*
	Optional Shared licenses: Participant or Server. For the Server, these licenses are available:[1]						*Optional Shared licenses: Participant or Server. For the Server, these licenses are available:*[1]					
	500-50,000 in increments of 500			*50,000-545,000 in increments of 1000*			*500-50,000 in increments of 500*			*50,000-545,000 in increments of 1000*		
	Optional FLEX license: 250						*Optional FLEX license: 250*					
IPSec VPN (sessions)[1]	250 (max. 250 combined IPSec and SSL VPN)						250 (max. 250 combined IPSec and SSL VPN)					
VPN Load Balancing	No support						Supported					
General Licenses												
Encryption	Base (DES)		*Opt. lic.: Strong (3DES/AES)*				Base (DES)		*Opt. lic.: Strong (3DES/AES)*			
Failover	No support						Active/Standby or Active/Active[1]					
Interface Speed	All: Fast Ethernet						Ethernet 0/0 and 0/1: Gigabit Ethernet[2] Ethernet 0/2, 0/3, and 0/4: Fast Ethernet					

Table 4-2 ASA 5510 Adaptive Security Appliance License Features (continued)

ASA 5510	Base License	Security Plus	
Security Contexts	No support	2	*Optional licenses:*
			5
VLANs, Maximum	50	100	

1. See the "License Notes" section.
2. Although the Ethernet 0/0 and 0/1 ports are Gigabit Ethernet, they are still identified as "Ethernet" in the software.

Table 4-3 ASA 5520 Adaptive Security Appliance License Features

ASA 5520	Base License							
Firewall Licenses								
Botnet Traffic Filter	Disabled		*Optional temporary license: Available*					
Firewall Conns, Concurrent	280 K							
GTP/GPRS	Disabled		*Optional license: Available*					
Unified Communications Proxy Sessions[1]	2	*Optional licenses:*						
		24	*50*	*100*	*250*	*500*	*750*	*1000*
VPN Licenses								
Adv. Endpoint Assessment	Disabled		*Optional license: Available*					
AnyConnect Essentials[1]	Disabled		*Optional license: Available*					
AnyConnect Mobile[1]	Disabled		*Optional license: Available*					
AnyConnect Premium SSL VPN (sessions)[1]	2	*Optional Permanent licenses:*						
		10	*25*	*50*	*100*	*250*	*500*	*750*
	Optional Shared licenses: Participant or Server. For the Server, these licenses are available:[1]							
	500-50,000 in increments of 500						*50,000-545,000 in increments of 1000*	
	Optional FLEX licenses:							
	250	*750*						
IPSec VPN (sessions)[1]	750 (max. 750 combined IPSec and SSL VPN)							
VPN Load Balancing	Supported							

Table 4-3 ASA 5520 Adaptive Security Appliance License Features (continued)

ASA 5520	Base License			
General Licenses				
Encryption	Base (DES)		*Optional license: Strong (3DES/AES)*	
Failover	Active/Standby or Active/Active[1]			
Security Contexts	2	*Optional licenses:*		
		5	*10*	*20*
VLANs, Maximum	150			

[1] See the "License Notes" section.

Table 4-4 ASA 5540 Adaptive Security Appliance License Features

ASA 5540	Base License									
Firewall Licenses										
Botnet Traffic Filter	Disabled		*Optional temporary license: Available*							
Firewall Conns, Concurrent	400 K									
GTP/GPRS	Disabled		*Optional license: Available*							
Unified Communications Proxy Sessions[1]	2	*Optional licenses:*								
		24	*50*	*100*	*250*	*500*	*750*	*1000*	*2000*	
VPN Licenses										
Adv. Endpoint Assessment	Disabled		*Optional license: Available*							
AnyConnect Essentials[1]	Disabled		*Optional license: Available*							
AnyConnect Mobile[1]	Disabled		*Optional license: Available*							
AnyConnect Premium SSL VPN (sessions)[1]	2	*Optional Permanent licenses:*								
		10	*25*	*50*	*100*	*250*	*500*	*750*	*1000*	*2500*
	Optional Shared licenses: Participant or Server. For the Server, these licenses are available:[1]									
	500-50,000 in increments of 500						*50,000-545,000 in increments of 1000*			
	Optional FLEX licenses:									
	250	*750*	*1000*	*2500*						

Table 4-4 ASA 5540 Adaptive Security Appliance License Features (continued)

ASA 5540	Base License				
IPSec VPN (sessions)[1]	5000 (max. 5000 combined IPSec and SSL VPN)				
VPN Load Balancing	Supported				
General Licenses					
Encryption	Base (DES)	*Optional license: Strong (3DES/AES)*			
Failover	Active/Standby or Active/Active[1]				
Security Contexts	2	*Optional licenses:*			
		5	*10*	*20*	*50*
VLANs, Maximum	200				

[1] See the “License Notes” section.

Table 4-5 ASA 5550 Adaptive Security Appliance License Features

ASA 5550	Base License										
Firewall Licenses											
Botnet Traffic Filter	Disabled	*Optional temporary license: Available*									
Firewall Conns, Concurrent	650 K										
GTP/GPRS	Disabled	*Optional license: Available*									
Unified Communications Proxy Sessions[1]	2	*Optional licenses:*									
		24	*50*	*100*	*250*	*500*	*750*	*1000*	*2000*	*3000*	
VPN Licenses											
Adv. Endpoint Assessment	Disabled	*Optional license: Available*									
AnyConnect Essentials[1]	Disabled	*Optional license: Available*									
AnyConnect Mobile[1]	Disabled	*Optional license: Available*									
AnyConnect Premium SSL VPN (sessions)[1]	2	*Optional Permanent licenses:*									
		10	*25*	*50*	*100*	*250*	*500*	*750*	*1000*	*2500*	*5000*
	Optional Shared licenses: Participant or Server. For the Server, these licenses are available:[1]										
	500-50,000 in increments of 500					*50,000-545,000 in increments of 1000*					
	Optional FLEX licenses:										
	250	*750*	*1000*	*2500*	*5000*						

Table 4-5 ASA 5550 Adaptive Security Appliance License Features (continued)

ASA 5550	Base License				
IPSec VPN (sessions)[1]	5000 (max. 5000 combined IPSec and SSL VPN)				
VPN Load Balancing	Supported				
General Licenses					
Encryption	Base (DES)	*Optional license: Strong (3DES/AES)*			
Failover	Active/Standby or Active/Active[1]				
Security Contexts	2	*Optional licenses:*			
		5	*10*	*20*	*50*
VLANs, Maximum	250				

[1] See the "License Notes" section.

Table 4-6 ASA 5580 Adaptive Security Appliance License Features

ASA 5580	Base License											
Firewall Licenses												
Botnet Traffic Filter	Disabled	*Optional temporary license: Available*										
Firewall Conns, Concurrent	650 K											
GTP/GPRS	Disabled	*Optional license: Available*										
Unified Communications Proxy Sessions[1]	2	*Optional licenses:*										
		24	*50*	*100*	*250*	*500*	*750*	*1000*	*2000*	*3000*	*5000*	*10000*[2]
VPN Licenses												
Adv. Endpoint Assessment	Disabled	*Optional license: Available*										
AnyConnect Essentials[1]	Disabled	*Optional license: Available*										
AnyConnect Mobile[1]	Disabled	*Optional license: Available*										
AnyConnect Premium SSL VPN (sessions)[1]	2	*Optional Permanent licenses:*										
		10	*25*	*50*	*100*	*250*	*500*	*750*	*1000*	*2500*	*5000*	
	Optional Shared licenses: Participant or Server. For the Server, these licenses are available:[1]											
	500-50,000 in increments of 500					*50,000-545,000 in increments of 1000*						
	Optional FLEX licenses:											
	250	*750*	*1000*	*2500*	*5000*							
IPSec VPN (sessions)[1]	5000 (max. 5000 combined IPSec and SSL VPN)											
VPN Load Balancing	Supported											
General Licenses												
Encryption	Base (DES)	*Optional license: Strong (3DES/AES)*										
Failover	Active/Standby or Active/Active[1]											

Table 4-6 ASA 5580 Adaptive Security Appliance License Features (continued)

ASA 5580	Base License				
Security Contexts	2	*Optional licenses:*			
		5	*10*	*20*	*50*
VLANs, Maximum	250				

1. See the "License Notes" section.
2. With the 10,000-session license, the total combined sessions can be 10,000, but the maximum number of Phone Proxy sessions is 5000.

License Notes

Table 4-7 lists footnotes for the tables in the "Licenses Per Model" section on page 4-1.

Table 4-7 License Notes

License	Notes
Active/Active failover	You cannot use Active/Active failover and VPN; if you want to use VPN, use Active/Standby failover.
AnyConnect Essentials	This license enables AnyConnect VPN client access to the adaptive security appliance. This license does not support deploy browser-based SSL VPN access or Cisco Secure Desktop. For these features, activate an AnyConnect Premium SSL VPN license instead of the AnyConnect Essentials license. **Note** With the AnyConnect Essentials license, VPN users can use a Web browser to log in, and download and start (WebLaunch) the AnyConnect client. The AnyConnect client software offers the same set of client features, whether it is enabled by this license or an AnyConnect Premium SSL VPN license. The AnyConnect Essentials license cannot be active at the same time as the following licenses on a given adaptive security appliance: AnyConnect Premium SSL VPN license (all types) or the Advanced Endpoint Assessment license. You can, however, run AnyConnect Essentials and AnyConnect Premium SSL VPN licenses on different adaptive security appliances in the same network. By default, the security appliance uses the AnyConnect Essentials license, but you can disable it to use other licenses by using the **no anyconnect-essentials** command.
AnyConnect Mobile	This license provides access to the AnyConnect Client for touch-screen mobile devices running Windows Mobile 5.0, 6.0, and 6.1. We recommend using this license if you want to support mobile access to AnyConnect 2.3 and later versions. This license requires activation of one of the following licenses to specify the total number SSL VPN sessions permitted: AnyConnect Essentials or AnyConnect Premium SSL VPN.
AnyConnect Premium SSL VPN Shared	A shared license lets the security appliance act as a shared license server for multiple client security appliances. The shared license pool is large, but the maximum number of sessions used by each individual security appliance cannot exceed the maximum number listed for permanent licenses.

Table 4-7 License Notes

License	Notes
Combined IPSec and SSL VPN sessions	• Although the maximum IPSec and SSL VPN sessions add up to more than the maximum VPN sessions, the combined sessions should not exceed the VPN session limit. If you exceed the maximum VPN sessions, you can overload the security appliance, so be sure to size your network appropriately. • If you start a clientless SSL VPN session and then start an AnyConnect client session from the portal, 1 session is used in total. However, if you start the AnyConnect client first (from a standalone client, for example) and then log into the clientless SSL VPN portal, then 2 sessions are used.
Unified Communications Proxy sessions	Phone Proxy, Mobility Advantage Proxy, Presence Federation Proxy, and TLS Proxy are all licensed under the UC Proxy umbrella, and can be mixed and matched. For example, if you configure a phone with a primary and backup Cisco Unified Communications Manager, there are 2 TLS/SRTP connections, so 2 UC Proxy sessions are used.

Information About Feature Licenses

A license specifies the options that are enabled on a given security appliance. It is represented by an activation key that is a 160-bit (5 32-bit words or 20 bytes) value. This value encodes the serial number (an 11 character string) and the enabled features.

This section includes the following topics:

- Preinstalled License, page 4-9
- Temporary, VPN Flex, and Evaluation Licenses, page 4-9
- Shared Licenses, page 4-12
- Licenses FAQ, page 4-16

Preinstalled License

By default, your security appliance ships with a license already installed. This license might be the Base License, to which you want to add more licenses, or it might already have all of your licenses installed, depending on what you ordered and what your vendor installed for you. See the "Viewing Your Current License" section on page 4-18 section to determine which licenses you have installed.

Temporary, VPN Flex, and Evaluation Licenses

In addition to permanent licenses, you can purchase a temporary license or receive an evaluation license that has a time-limit. For example, you might buy a VPN Flex license to handle short-term surges in the number of concurrent SSL VPN users, or you might order a Botnet Traffic Filter temporary license that is valid for 1 year.

This section includes the following topics:

- How the Temporary License Timer Works, page 4-10

- How Multiple Licenses Interact, page 4-10
- Failover and Temporary Licenses, page 4-11

How the Temporary License Timer Works

- The timer for the temporary license starts counting down when you activate it on the security appliance.
- If you stop using the temporary license before it times out, for example you activate a permanent license or a different temporary license, then the timer halts. The timer only starts again when you reactivate the temporary license.
- If the temporary license is active, and you shut down the security appliance, then the timer continues to count down. If you intend to leave the security appliance in a shut down state for an extended period of time, then you should activate the permanent license before you shut down to preserve the temporary license.
- When a temporary license expires, the next time you reload the security appliance, the permanent license is used; you are not forced to perform a reload immediately when the license expires.

Note We suggest you do not change the system clock after you install the temporary license. If you set the clock to be a later date, then if you reload, the security appliance checks the system clock against the original installation time, and assumes that more time has passed than has actually been used. If you set the clock back, and the actual running time is greater than the time between the original installation time and the system clock, then the license immediately expires after a reload.

How Multiple Licenses Interact

- When you activate a temporary license, then features from both permanent and temporary licenses are merged to form the running license. Note that the security appliance only uses the *highest* value from each license for each feature; the values are not added together. The security appliance displays any resolved conflicts between the licenses when you enter a temporary activation key. In the rare circumstance that a temporary license has lower capability than the permanent license, the permanent license values are used.
- When you activate a permanent license, it overwrites the currently-running permanent and temporary licenses and becomes the running license.

  Note If you install a new permanent license, and it is a downgrade from the temporary license, then you need to reload the security appliance to disable the temporary license and restore the permanent license. Until you reload, the temporary license continues to count down.

  If you reactivate the *already installed* permanent license, you do not need to reload the security appliance; the temporary license does not continue to count down, and there is no disruption of traffic.

- To reenable the features of the temporary license if you later activate a permanent license, simply reenter the temporary activation key. For a license upgrade, you do not need to reload.

- To switch to a different temporary license, enter the new activation key; the new license is used instead of the old temporary license and combines with the permanent license to create a new running license. The security appliance can have multiple temporary licenses installed; but only one is active at any given time.

See the following figure for examples of permanent and VPN Flex activation keys, and how they interact.

Figure 4-1 Permanent and VPN Flex Activation Keys

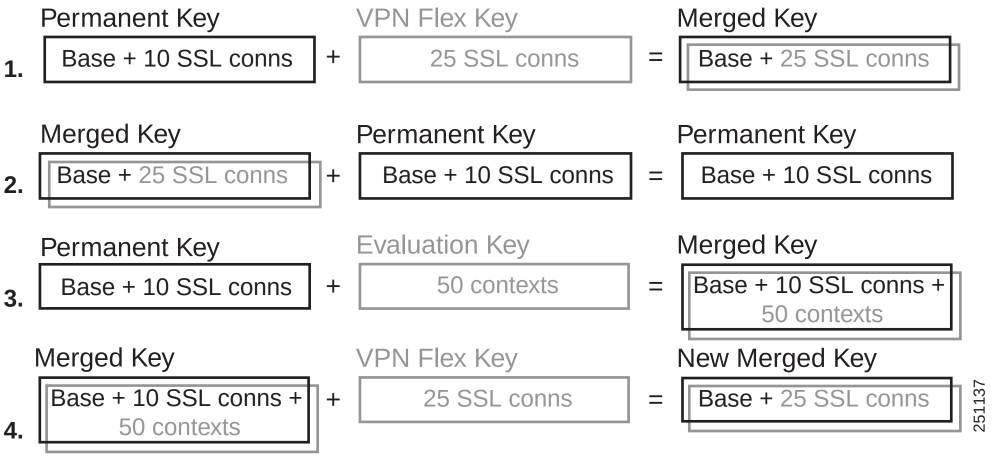

1. In example 1 in the above figure, you apply a temporary key with 25 SSL sessions; because the VPN Flex value is greater than the permanent key value of 10 sessions, the resulting running key is a merged key that uses the VPN Flex value of 25 sessions, and not a combined total of 35 sessions.
2. In example 2 above, the merged key from example 1 is replaced by the permanent key, and the VPN Flex license is disabled. The running key defaults to the permanent key value of 10 sessions.
3. In example 3 above, an evaluation license including 50 contexts is applied to the permanent key, so the resulting running key is a merged key that includes all the features of the permanent key plus the 50 context license.
4. In example 4 above, the merged key from example 3 has the VPN Flex key applied. Because the security appliance can only use one temporary key at a time, the VPN flex key replaces the evaluation key, so the end result is the same as the merged key from example 1.

Failover and Temporary Licenses

With failover, identical licenses are required. For failover purposes, temporary and permanent licenses appear to be identical, so you can have a permanent license on one unit and a temporary license on the other unit. This functionality is useful in an emergency situation; for example, if one of your units fails, and you have an extra unit, you can install the extra unit while the other one is repaired. If you do not normally use the extra unit for SSL VPN, then a VPN Flex license is a perfect solution while the other unit is being repaired.

Because the temporary license continues to count down for as long as it is activated on a failover unit, we do not recommend using a temporary license in a permanent failover installation; when the temporary license expires, failover will no longer work.

Shared Licenses

A shared license lets you purchase a large number of SSL VPN sessions and share the sessions as needed amongst a group of security appliances by configuring one of the security appliances as a shared licensing server, and the rest as shared licensing participants. This section describes how a shared license works, and includes the following topics:

- Information About the Shared Licensing Server and Participants, page 4-12
- Communication Issues Between Participant and Server, page 4-13
- Information About the Shared Licensing Backup Server, page 4-13
- Failover and Shared Licenses, page 4-14
- Maximum Number of Participants, page 4-15

Information About the Shared Licensing Server and Participants

The following steps describe how shared licenses operate:

1. Decide which security appliance should be the shared licensing server, and purchase the shared licensing server license using that device serial number.
2. Decide which security appliances should be shared licensing participants, including the shared licensing backup server, and obtain a shared licensing participant license for each device, using each device serial number.
3. (Optional) Designate a second security appliance as a shared licensing backup server. You can only specify one backup server.

   Note The shared licensing backup server only needs a participant license.

4. Configure a shared secret on the shared licensing server; any participants with the shared secret can use the shared license.
5. When you configure the security appliance as a participant, it registers with the shared licensing server by sending information about itself, including the local license and model information.

   Note The participant needs to be able to communicate with the server over the IP network; it does not have to be on the same subnet.

6. The shared licensing server responds with information about how often the participant should poll the server.
7. When a participant uses up the sessions of the local license, it sends a request to the shared licensing server for additional sessions in 50-session increments.
8. The shared licensing server responds with a shared license. The total sessions used by a participant cannot exceed the maximum sessions for the platform model.

   Note The shared licensing server can also participate in the shared license pool. It does not need a participant license as well as the server license to participate.

   a. If there are not enough sessions left in the shared license pool for the participant, then the server responds with as many sessions as available.
   b. The participant continues to send refresh messages requesting more sessions until the server can adequately fulfill the request.
9. When the load is reduced on a participant, it sends a message to the server to release the shared sessions.

Note The security appliance uses SSL between the server and participant to encrypt all communications.

Communication Issues Between Participant and Server

See the following guidelines for communication issues between the participant and server:

- If a participant fails to send a refresh after 3 times the refresh interval, then the server releases the sessions back into the shared license pool.
- If the participant cannot reach the license server to send the refresh, then the participant can continue to use the shared license it received from the server for up to 24 hours.
- If the participant is still not able to communicate with a license server after 24 hours, then the participant releases the shared license, even if it still needs the sessions. The participant leaves existing connections established, but cannot accept new connections beyond the license limit.
- If a participant reconnects with the server before 24 hours expires, but after the server expired the participant sessions, then the participant needs to send a new request for the sessions; the server responds with as many sessions as can be reassigned to that participant.

Information About the Shared Licensing Backup Server

The shared licensing backup server must register successfully with the main shared licensing server before it can take on the backup role. When it registers, the main shared licensing server syncs server settings as well as the shared license information with the backup, including a list of registered participants and the current license usage. The main server and backup server sync the data at 10 second intervals. After the initial sync, the backup server can successfully perform backup duties, even after a reload.

When the main server goes down, the backup server takes over server operation. The backup server can operate for up to 30 continuous days, after which the backup server stops issuing sessions to participants, and existing sessions time out. Be sure to reinstate the main server within that 30-day period. Critical-level syslog messages are sent at 15 days, and again at 30 days.

When the main server comes back up, it syncs with the backup server, and then takes over server operation.

When the backup server is not active, it acts as a regular participant of the main shared licensing server.

Note When you first launch the main shared licensing server, the backup server can only operate independently for 5 days. The operational limit increases day-by-day, until 30 days is reached. Also, if the main server later goes down for any length of time, the backup server operational limit decrements day-by-day. When the main server comes back up, the backup server starts to increment again day-by-day. For example, if the main server is down for 20 days, with the backup server active during that time, then the backup server will only have a 10-day limit left over. The backup server "recharges" up to the maximum 30 days after 20 more days as an inactive backup. This recharging function is implemented to discourage misuse of the shared license.

Failover and Shared Licenses

This section describes how shared licenses interact with failover, and includes the following topics:

- "Failover and Shared License Servers" section on page 4-14
- "Failover and Shared License Participants" section on page 4-15

Failover and Shared License Servers

This section describes how the main server and backup server interact with failover. Because the shared licensing server is also performing normal duties as the security appliance, including performing functions such as being a VPN gateway and firewall, then you might need to configure failover for the main and backup shared licensing servers for increased reliability.

Note The backup server mechanism is separate from, but compatible with, failover.

Shared licenses are supported only in single context mode, so Active/Active failover is not supported.

Both main shared licensing server units in the failover pair need to have the same license. So if you purchase a 10,000 session shared license for the primary main server unit, you must also purchase a 10,000 session shared license for the standby main server unit. Because the standby unit does not pass traffic when it is in a standby state, the total number of sessions remains at 10,000 in this example, *not* a combined 20,000 sessions.

For Active/Standby failover, the primary unit acts as the main shared licensing server, and the standby unit acts as the main shared licensing server after failover; because both units need to have the same license, both units can act as the main licensing server. The standby unit does *not* act as the backup shared licensing server. Instead, you can have a second pair of units acting as the backup server, if desired.

For example, you have a network with 2 failover pairs. Pair #1 includes the main licensing server. Pair #2 includes the backup server. When the primary unit from Pair #1 goes down, the standby unit immediately becomes the new main licensing server. The backup server from Pair #2 never gets used. Only if both units in Pair #1 go down does the backup server in Pair #2 come into use as the shared licensing server. If Pair #1 remains down, and the primary unit in Pair #2 goes down, then the standby unit in Pair #2 comes into use as the shared licensing server (see Figure 4-2).

Figure 4-2 Failover and Shared License Servers

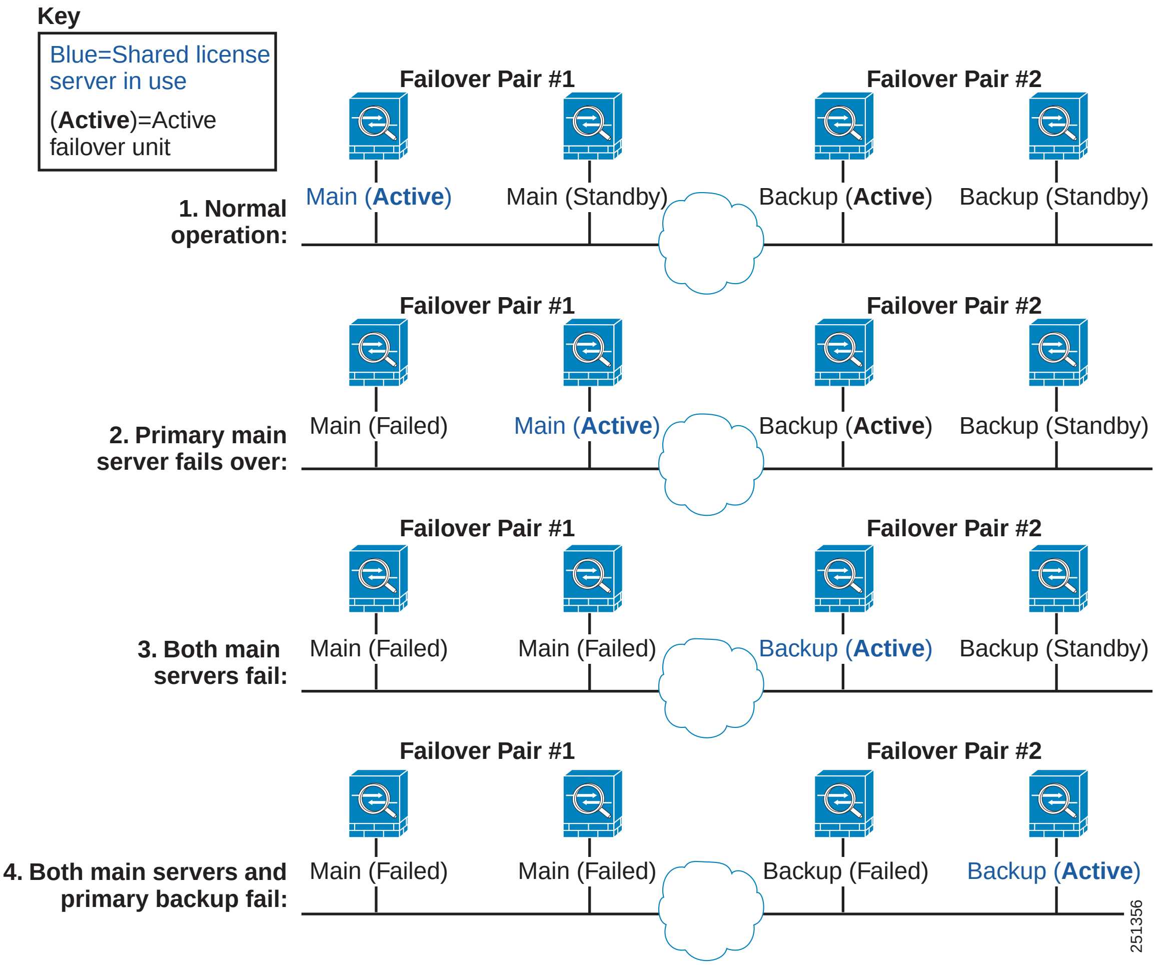

The standby backup server shares the same operating limits as the primary backup server; if the standby unit becomes active, it continues counting down where the primary unit left off. See the "Information About the Shared Licensing Backup Server" section on page 4-13 for more information.

Failover and Shared License Participants

For participant pairs, both units register with the shared licensing server using separate participant IDs. The active unit syncs its participant ID with the standby unit. The standby unit uses this ID to generate a transfer request when it switches to the active role. This transfer request is used to move the shared sessions from the previously active unit to the new active unit.

Maximum Number of Participants

The security appliance does not limit the number of participants for the shared license; however, a very large shared network could potentially affect the performance on the licensing server. In this case, you can increase the delay between participant refreshes, or you can create two shared networks.

Licenses FAQ

Q. Can I activate multiple temporary licenses, for example, VPN Flex and Botnet Traffic Filter?

A. No. You can only use one temporary license at a time. The last license you activate is the one in use. In the case of evaluation licenses that group multiple features into one activation key, then multiple features are supported at the same time. But temporary licenses for sale by Cisco are limited to one feature per activation key.

Q. Can I “stack” temporary licenses so that when the time limit runs out, it will automatically use the next license?

A. No. You can install multiple temporary licenses, but only the last activated license is active. When the active license expires, you need to manually activate the new one. Be sure to activate it shortly *before* the old one expires so you do not lose functionality. (Any remaining time on the old license remains unused; for example, if you use 10 months of a 12-month license, and activate a new 12-month license, then the remaining 2 months of the first license goes unused unless you later reactivate it. We recommend that you activate the new license as close as possible to the end of the old license to maximize the license usage.)

Q. Can I install a new permanent license while maintaining an active temporary license?

A. No. The temporary license will be deactivated when you apply a permanent license. You have to activate the permanent license, and then reactivate the temporary license to be able to use the new permanent license along with the temporary license. This will cause temporary loss of functionality for the features reliant on the temporary license.

Q. For failover, can I use a shared licensing server as the primary unit, and the shared licensing backup server as the secondary unit?

A. No. The secondary unit must also have a shared licensing server license. The backup server, which has a participant license, can be in a separate failover pair of two backup servers.

Q. Do I need to buy the same licenses for the secondary unit in a failover pair? Even for a shared licensing server?

A. Yes. Both units need the same licenses. For a shared licensing server, you need to buy the same shared licensing server license for both units. **Note:** In Active/Standby failover, for licenses that specify the number of sessions, the sessions for both units are not added to each other; only the active unit sessions can be used. For example, for a shared SSL VPN license, you need to purchase a 10,000 user session for both the active and the standby unit; the total number of sessions is 10,000, *not* 20,000 combined.

Q. Can I use a VPN Flex or permanent SSL VPN license in addition to a shared SSL VPN license?

A. Yes. The shared license is used only after the sessions from the locally installed license (VPN Flex or permanent) are used up. **Note**: On the shared licensing server, the permanent SSL VPN license is not used; you can however use a VPN Flex license at the same time as the shared licensing server license. In this case, the VPN Flex license sessions are available for local SSL VPN sessions only; they cannot be added to the shared licensing pool for use by participants.

Guidelines and Limitations

See the following guidelines for activation keys.

Context Mode Guidelines

- In multiple context mode, apply the activation key in the system execution space.
- Shared licenses are not supported in multiple context mode.

Firewall Mode Guidelines

All license types are available in both routed and transparent mode.

Failover Guidelines

- You must have the same licenses activated on the primary and secondary units.

  Note For failover purposes, there is no distinction between permanent and temporary licenses as long as the feature set is the same between the two units. See the "Failover and Temporary Licenses" section on page 4-11 for more information.

- Shared licenses are not supported in Active/Active mode. See the "Failover and Shared Licenses" section on page 4-14 for more information.

Upgrade Guidelines

Your activation key remains compatible if you upgrade to Version 8.2 or later, and also if you later downgrade. After you upgrade, if you activate additional feature licenses that were introduced *before 8.2*, then the activation key continues to be compatible with earlier versions if you downgrade. However if you activate feature licenses that were introduced in *8.2 or later*, then the activation key is not backwards compatible. If you have an incompatible license key, then see the following guidelines:

- If you previously entered an activation key in an earlier version, then the security appliance uses that key (without any of the new licenses you activated in Version 8.2 or later).
- If you have a new system and do not have an earlier activation key, then you need to request a new activation key compatible with the earlier version.

Additional Guidelines and Limitations

- The activation key is not stored in your configuration file; it is stored as a hidden file in Flash memory.
- The activation key is tied to the serial number of the device. Feature licenses cannot be transferred between devices (except in the case of a hardware failure). If you have to replace your device due to a hardware failure, contact the Cisco Licensing Team to have your existing license transferred to the new serial number. The Cisco Licensing Team will ask for the Product Authorization Key reference number and existing serial number.
- Once purchased, you cannot return a license for a refund or for an upgraded license.
- You cannot add two separate licenses for the same feature together; for example, if you purchase a 25-session SSL VPN license, and later purchase a 50-session license, you cannot use 75 sessions; you can use a maximum of 50 sessions.
- Although you can activate all license types, some features are incompatible with each other; for example, multiple context mode and VPN. In the case of the AnyConnect Essentials license, the license is incompatible with the following licenses: full SSL VPN license, shared SSL VPN license,

and Advanced Endpoint Assessment license. By default, the AnyConnect Essentials license is used instead of the above licenses, but you can disable the AnyConnect Essentials license in the configuration to restore use of the other licenses using the **no anyconnect-essentials** command.

Viewing Your Current License

This section describes how to view your current license, and for temporary activation keys, how much time the license has left.

To view the current license, choose Configuration > Device Management > Licensing > Activation Key.

In multiple context mode, view the activation key in the System execution space by choosing Configuration > Device Management > Activation Key.

Obtaining an Activation Key

To obtain an activation key, you need a Product Authorization Key, which you can purchase from your Cisco account representative. You need to purchase a separate Product Activation Key for each feature license. For example, if you have the Base License, you can purchase separate keys for Advanced Endpoint Assessment and for additional SSL VPN sessions.

Note For a failover pair, you need separate activation keys for each unit. Make sure the licenses included in the keys are the same for both units.

After obtaining the Product Authorization Keys, register them on Cisco.com by performing the following steps:

Step 1 Obtain the serial number for your security appliance by choosing Configuration > Device Management > Licensing > Activation Key (in multiple context mode, view the serial number in the System execution space).

Step 2 Access one of the following URLs.

- Use the following website if you are a registered user of Cisco.com:

  `http://www.cisco.com/go/license`

- Use the following website if you are not a registered user of Cisco.com:

  `http://www.cisco.com/go/license/public`

Step 3 Enter the following information, when prompted:

- Product Authorization Key (if you have multiple keys, enter one of the keys first. You have to enter each key as a separate process.)
- The serial number of your security appliance
- Your email address

An activation key is automatically generated and sent to the email address that you provide. This key includes all features you have registered so far for permanent licenses. For VPN Flex licenses, each license has a separate activation key.

Step 4 If you have additional Product Authorization Keys, repeat Step 3 for each Product Authorization Key. After you enter all of the Product Authorization Keys, the final activation key provided includes all of the permanent features you registered.

Entering a New Activation Key

This section describes how to enter a new activation key.

Prerequisites

- Before entering the activation key, ensure that the image in Flash memory and the running image are the same by entering the show activation-key command. You can do this by reloading the security appliance before entering the new activation key.
- If you are already in multiple context mode, enter the activation key in the system execution space.
- Some licenses require you to reload the security appliance after you activate them. Table 4-8 lists the licenses that require reloading.

Table 4-8 License Reloading Requirements

Model	License Action Requiring Reload
ASA 5505 and ASA 5510	Changing between the Base and Security Plus license.
All models	Changing the Encryption license.
All models	Downgrading any license (for example, going from 10 contexts to 2 contexts). **Note** If a temporary license expires, and the permanent license is a downgrade, then you do not need to immediately reload the security appliance; the next time you reload, the permanent license is restored.

Limitations and Restrictions

Your activation key remains compatible if you upgrade to Version 8.2 or later, and also if you later downgrade. After you upgrade, if you activate additional feature licenses that were introduced *before 8.2*, then the activation key continues to be compatible with earlier versions if you downgrade. However if you activate feature licenses that were introduced in *8.2 or later*, then the activation key is not backwards compatible. If you have an incompatible license key, then see the following guidelines:

- If you previously entered an activation key in an earlier version, then the security appliance uses that key (without any of the new licenses you activated in Version 8.2 or later).
- If you have a new system and do not have an earlier activation key, then you need to request a new activation key compatible with the earlier version.

Detailed Steps

Step 1 Choose Configuration > Device Management > Licensing > Activation Key.

Step 2 Enter the new activation key in the New Activation Key field.

The key is a five-element hexadecimal string with one space between each element. The leading 0x specifier is optional; all values are assumed to be hexadecimal. For example:

```
0xd11b3d48 0xa80a4c0a 0x48e0fd1c 0xb0443480 0x843fc490
```

You can enter one permanent key, and multiple temporary keys. The last temporary key entered is the active one. See the "Temporary, VPN Flex, and Evaluation Licenses" section on page 4-9 for more information. To change the running activation key, enter a new value.

Step 3 Click **Update Activation Key**.

Upgrading the License for a Failover Pair

If you need to upgrade the license on a failover pair, you might have some amount of downtime depending on whether the license requires a reload. See Table 4-8 on page 4-19 for more information about licenses requiring a reload. This section includes the following topics:

- Upgrading the License for a Failover (No Reload Required), page 4-20
- Upgrading the License for a Failover (Reload Required), page 4-21

Upgrading the License for a Failover (No Reload Required)

Use the following procedure if your new license does not require you to reload. See Table 4-8 on page 4-19 for more information about licenses requiring a reload. This procedure ensures that there is no downtime.

Prerequisites

Before you upgrade the license, be sure that both units are operating correctly, the Failover LAN interface is up, and there is not an imminent failover event; for example, monitored interfaces are operating normally.

On each unit, go to Monitoring > Properties > Failover > Status to view the failover status and the monitored interface status.

Detailed Steps

Step 1 On the active unit, choose Configuration > Device Management > High Availability > Failover > Setup, and uncheck the **Enable Failover** check box.

The standby unit remains in a pseudo-standby state. Deactivating failover on the active unit prevents the standby unit from attempting to become active during the period when the licenses do not match.

Step 2 Click **Apply**.

Step 3 Choose Configuration > Device Management > Licensing > Activation Key, and enter the new activation key that you obtained with the active unit serial number.

Step 4 Click **Update Activation Key**.

Step 5 Log into the standby unit by double-clicking its address in the Device List.

If you the device is not in the Device List, click **Add** to add the device. You might be prompted for credentials to log in.

Step 6 Choose Configuration > Device Management > Licensing > Activation Key, and enter the new activation key that you obtained with the standby unit serial number.

Step 7 Click **Update Activation Key**.

Step 8 Log into the active unit again by double-clicking its address in the Device List.

Step 9 Choose Configuration > Device Management > High Availability > Failover > Setup, and re-check the **Enable Failover** check box.

Step 10 Click **Apply**.

Upgrading the License for a Failover (Reload Required)

Use the following procedure if your new license requires you to reload. See Table 4-8 on page 4-19 for more information about licenses requiring a reload. Reloading the failover pair causes a loss of connectivity during the reload.

Prerequisites

Before you upgrade the license, be sure that both units are operating correctly, the Failover LAN interface is up, and there is not an imminent failover event; for example, monitored interfaces are operating normally.

On each unit, choose Monitoring > Properties > Failover > Status to view the failover status and the monitored interface status.

Detailed Steps

Step 1 On the active unit, choose Configuration > Device Management > High Availability > Failover > Setup, and uncheck the **Enable Failover** check box.

The standby unit remains in a pseudo-standby state. Deactivating failover on the active unit prevents the standby unit from attempting to become active during the period when the licenses do not match.

Step 2 Click **Apply**.

Step 3 Choose Configuration > Device Management > Licensing > Activation Key, and enter the new activation key that you obtained with the active unit serial number.

Step 4 Click **Update Activation Key**.

Step 5 Log into the standby unit by double-clicking its address in the Device List.

If you the device is not in the Device List, click **Add** to add the device. You might be prompted for credentials to log in.

Step 6 Choose Configuration > Device Management > Licensing > Activation Key, and enter the new activation key that you obtained with the standby unit serial number.

Step 7 Click **Update Activation Key**.

Step 8 Log into the active unit again by double-clicking its address in the Device List.

Step 9 Choose Configuration > Device Management > High Availability > Failover > Setup, and recheck the **Enable Failover** check box.

Step 10 Click **Apply**.

Step 11 Schedule a reload of the security appliance by choosing Tools > System Reload.

Step 12 Choose the reload options to reload the security appliance at a time you desire, and click **Schedule Reload**.

Choose a time when the loss of service has the least impact.

Step 13 Log into the standby unit again by double-clicking its address in the Device List.

Step 14 Schedule a reload of the security appliance by choosing Tools > System Reload.

Step 15 Choose the reload options to reload the security appliance at the same time you choose for the active unit, and click **Schedule Reload**.

Both units will reload at the same time, and the new licenses will be in effect.

Configuring a Shared License

This section describes how to configure the shared licensing server and participants. For more information about shared licenses, see the "Shared Licenses" section on page 4-12.

This section includes the following topics:

- Configuring the Shared Licensing Server, page 4-22
- Configuring the Shared Licensing Participant and the Optional Backup Server, page 4-23
- Monitoring the Shared License, page 4-24

Configuring the Shared Licensing Server

This section describes how to configure the security appliance to be a shared licensing server.

Prerequisites

The server must have a shared licensing server key.

Detailed Steps

Step 1 Choose **Configuration > Device Management > Licenses > Shared SSL VPN Licenses**.

Step 2 In the Shared Secret field, enter the shared secret as a string between 4 and 128 ASCII characters.

Any participant with this secret can use the license server.

Step 3 (Optional) In the TCP IP Port field, enter the port on which the server listens for SSL connections from participants, between 1 and 65535.

The default is TCP port 50554.

Step 4 (Optional) In the Refresh interval field, enter the refresh interval between 10 and 300 seconds.

This value is provided to participants to set how often they should communicate with the server. The default is 30 seconds.

Step 5 In the Interfaces that serve shared licenses area, check the **Shares Licenses** check box for any interfaces on which participants contact the server.

Step 6 (Optional) To identify a backup server, in the Optional backup shared SSL VPN license server area:

a. In the Backup server IP address field, enter the backup server IP address.

b. In the Primary backup server serial number field, enter the backup server serial number.

c. If the backup server is part of a failover pair, identify the standby unit serial number in the Secondary backup server serial number field.

You can only identify 1 backup server and its optional standby unit.

Step 7 Click **Apply**.

What to Do Next

See the“Configuring the Shared Licensing Participant and the Optional Backup Server” section on page 4-23.

Configuring the Shared Licensing Participant and the Optional Backup Server

This section configures a shared licensing participant to communicate with the shared licensing server; this section also describes how you can optionally configure the participant as the backup server.

Prerequisites

The participant must have a shared licensing participant key.

Detailed Steps

Step 1 Choose the Configuration > Device Management > Licenses > Shared SSL VPN Licenses pane.

Step 2 In the Shared Secret field, enter the shared secret as a string between 4 and 128 ASCII characters.

Step 3 (Optional) In the TCP IP Port field, enter the port on which to communicate with the server using SSL, between 1 and 65535.

The default is TCP port 50554.

Step 4 (Optional) To identify the participant as the backup server, in the Select backup role of participant area:

a. Click the **Backup Server** radio button.

b. Check the **Shares Licenses** check box for any interfaces on which participants contact the backup server.

Step 5 Click **Apply**.

Monitoring the Shared License

To monitor the shared license, choose Monitoring > VPN > Clientless SSL VPN > Shared Licenses.

Feature History for Licensing

Table 4-9 lists the release history for this feature.

Table 4-9 Feature History for Licensing

Feature Name	Releases	Feature Information
Increased Connections and VLANs	7.0(5)	Increased the following limits: • ASA5510 Base license connections from 32000 to 5000; VLANs from 0 to 10. • ASA5510 Security Plus license connections from 64000 to 130000; VLANs from 10 to 25. • ASA5520 connections from 130000 to 280000; VLANs from 25 to 100. • ASA5540 connections from 280000 to 400000; VLANs from 100 to 200.
SSL VPN Licenses	7.1(1)	SSL VPN licenses were introduced.
Increased SSL VPN Licenses	7.2(1)	A 5000-user SSL VPN license was introduced for the ASA 5550 and above.
Increased VLANs	7.2(2)	The maximum number of VLANs for the Security Plus license on the ASA 5505 security appliance was increased from 5 (3 fully functional; 1 failover; one restricted to a backup interface) to 20 fully functional interfaces. In addition, the number of trunk ports was increased from 1 to 8. Now there are 20 fully functional interfaces, you do not need to use the backup interface command to cripple a backup ISP interface; you can use a fully-functional interface for it. The backup interface command is still useful for an Easy VPN configuration. VLAN limits were also increased for the ASA 5510 security appliance (from 10 to 50 for the Base license, and from 25 to 100 for the Security Plus license), the ASA 5520 adaptive security appliance (from 100 to 150), the ASA 5550 adaptive security appliance (from 200 to 250).

Table 4-9 Feature History for Licensing (continued)

Feature Name	Releases	Feature Information
Gigabit Ethernet Support for the ASA 5510 Security Plus License	7.2(3)	The ASA 5510 security appliance now supports Gigabit Ethernet (1000 Mbps) for the Ethernet 0/0 and 0/1 ports with the Security Plus license. In the Base license, they continue to be used as Fast Ethernet (100 Mbps) ports. Ethernet 0/2, 0/3, and 0/4 remain as Fast Ethernet ports for both licenses. **Note** The interface names remain Ethernet 0/0 and Ethernet 0/1.
Advanced Endpoint Assessment License	8.0(2)	The Advanced Endpoint Assessment license was introduced. As a condition for the completion of a Cisco AnyConnect or clientless SSL VPN connections, the remote computer scans for a greatly expanded collection of antivirus and antispyware applications, firewalls, operating systems, and associated updates. It also scans for any registry entries, filenames, and process names that you specify. It sends the scan results to the adaptive security appliance. The security appliance uses both the user login credentials and the computer scan results to assign a Dynamic Access Policy (DAP). With an Advanced Endpoint Assessment License, you can enhance Host Scan by configuring an attempt to update noncompliant computers to meet version requirements. Cisco can provide timely updates to the list of applications and versions that Host Scan supports in a package that is separate from Cisco Secure Desktop.
VPN Load Balancing for the ASA 5510	8.0(2)	VPN load balancing is now supported on the ASA 5510 Security Plus license.
AnyConnect for Mobile License	8.0(3)	The AnyConnect for Mobile license lets Windows mobile devices connect to the security appliance using the AnyConnect client.
VPN Flex and Evaluation Licenses	8.0(4)/8.1(2)	Support for temporary licenses was introduced. VPN Flex licenses provide temporary support for extra SSL VPN sessions.
Increased VLANs for the ASA 5580	8.1(2)	The number of VLANs supported on the ASA 5580 are increased from 100 to 250.
Unified Communications Proxy Sessions license	8.0(4)	The UC Proxy sessions license was introduced. This feature is not available in Version 8.1.
Botnet Traffic Filter License	8.2(1)	The Botnet Traffic Filter license was introduced. The Botnet Traffic Filter protects against malware network activity by tracking connections to known bad domains and IP addresses.

Table 4-9 Feature History for Licensing (continued)

Feature Name	Releases	Feature Information
AnyConnect Essentials License	8.2(1)	This license enables AnyConnect VPN client access to the adaptive security appliance. This license does not support browser-based SSL VPN access or Cisco Secure Desktop. For these features, activate an AnyConnect Premium SSL VPN license instead of the AnyConnect Essentials license. **Note** With the AnyConnect Essentials license, VPN users can use a Web browser to log in, and download and start (WebLaunch) the AnyConnect client. The AnyConnect client software offers the same set of client features, whether it is enabled by this license or an AnyConnect Premium SSL VPN license. The AnyConnect Essentials license cannot be active at the same time as the following licenses on a given adaptive security appliance: AnyConnect Premium SSL VPN license (all types) or the Advanced Endpoint Assessment license. You can, however, run AnyConnect Essentials and AnyConnect Premium SSL VPN licenses on different adaptive security appliances in the same network. By default, the security appliance uses the AnyConnect Essentials license, but you can disable it to use other licenses by using the **no anyconnect-essentials** command.
Shared Licenses for SSL VPN	8.2(1)	Shared licenses for SSL VPN were introduced. Multiple security appliances can share a pool of SSL VPN sessions on an as-needed basis.

CHAPTER 5

Defining Preferences and Using Configuration, Diagnostic, and File Management Tools

This chapter describes the preferences and tools available for configuration, problem diagnosis, and file management, and includes the following sections:

Defining Preferences

This feature lets you change the behavior of some ASDM functions between sessions.

To change various settings in ASDM, perform the following steps:

Step 1 In the main ASDM application window, choose **Tools** > **Preferences**.

The Preferences dialog box appears, with three tabs: General, Rules Table, and Syslog.

Step 2 To define your settings, click one of these tabs: the **General** tab to specify general preferences; the **Rules Tables** tab to specify preferences for the Rules table; and the **Syslog** tab to specify the appearance of syslog messages displayed in the Home pane and to enable the display of a warning message for NetFlow-related syslog messages.

Step 3 On the General tab, specify the following:

a. Check the **Warn that configuration in ASDM is out of sync with the configuration in ASA** check box to be notified when the startup configuration and the running configuration are no longer in sync with each other.

b. Check the **Show configuration restriction message to read-only user** check box to display the following message to a read-only user at startup. This option is checked by default.

```
"You are not allowed to modify the ASA configuration, because you do not have
sufficient privileges."
```

c. Check the **Confirm before exiting ASDM** check box to display a prompt when you try to close ASDM to confirm that you want to exit. This option is checked by default.

d. Check the **Enable screen reader support (requires ASDM restart)** check box to enable screen readers to work. You must restart ASDM to enable this option.

e. Check the **Preview commands before sending them to the device** check box to view CLI commands generated by ASDM.

f. Check the **Enable cumulative (batch) CLI delivery** check box to send multiple commands in a single group to the security appliance.

g. Enter the minimum amount of time in seconds for a configuration to send a timeout message. The default is 60 seconds.

h. To allow the Packet Capture Wizard to display captured packets, enter the name of the network sniffer application or click **Browse** to find it in the file system.

Step 4 On the Rules Tables tab, specify the following:

a. Display settings let you change the way rules appear in the Rules table.

- Check the **Auto-expand network and service object groups with specified prefix** check box to display the network and service object groups automatically expanded based on the Auto-Expand Prefix setting.
- In the Auto-Expand Prefix field, enter the prefix of the network and service object groups to expand automatically when displayed.
- Check the **Show members of network and service object groups** check box to display members of network and service object groups and the group name in the Rules table. If the check box is not checked, only the group name is displayed.
- In the Limit Members To field, enter the number of network and service object groups to display. When the object group members are displayed, then only the first *n* members are displayed.
- Check the **Show all actions for service policy rules** check box to display all actions in the Rules table. When unchecked, a summary appears.

b. Deployment settings let you configure the behavior of the security appliance when deploying changes to the Rules table.

- Check the **Issue "clear xlate" command when deploying access lists** check box to clear the NAT table when deploying new access lists. This setting ensures the access lists that are configured on the security appliance are applied to all translated addresses.

c. Access Rule Hit Count Settings let you configure the frequency for which the hit counts are updated in the Access Rules table. Hit counts are applicable for explicit rules only. No hit count will be displayed for implicit rules in the Access Rules table.

- Check the **Update access rule hit counts automatically** check box to have the hit counts automatically updated in the Access Rules table.
- In the Update Frequency field, specify the frequency in seconds in which the hit count column is updated in the Access Rules table. Valid values are 10 - 86400 seconds.

Step 5 On the Syslog tab, specify the following:

- In the Syslog Colors area, you can customize the message display by configuring background or foreground colors for messages at each severity level. The Severity column lists each severity level by name and number. To change the background color or foreground color for messages at a specified severity level, click the corresponding column. The Pick a Color dialog box appears. Click one of the following tabs:
  - On the Swatches tab, choose a color from the palette, and click **OK**.
  - On the HSB tab, specify the H, S, and B settings, and click **OK**.
  - On the RGB tab, specify the Red, Green, and Blue settings, and click **OK**.

- In the NetFlow area, to enable the display of a warning message to disable redundant syslog messages, check the **Warn to disable redundant syslog messages when NetFlow action is first applied to the global service policy rule** check box.

Step 6 After you have specified settings on these three tabs, click **OK** to save your settings and close the Preferences dialog box.

Note Each time that you check or uncheck a preferences setting, the change is saved to the .conf file and becomes available to all the other ASDM sessions running on the workstation at the time. You must restart ASDM for all changes to take effect.

Modes

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

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

Using Configuration Tools

This section includes the following topics:

- Resetting a Device to the Factory Default Configuration, page 5-3
- Save the Running Configuration to a TFTP Server, page 5-4
- Saving an Internal Log Buffer to Flash, page 5-5
- Using the Command Line Interface, page 5-5
- Showing Commands Ignored by ASDM on the Device, page 5-6

Resetting a Device to the Factory Default Configuration

The default configuration provides the minimum commands required to connect to the security appliance using ASDM.

Note This feature is available only in routed firewall mode; transparent mode does not support IP addresses for interfaces. In addition, this feature is available only in single context mode; an security appliance with a cleared configuration does not have any defined contexts to configure automatically using this feature.

To reset the adaptive security appliance to the factory default configuration, perform the following steps:

Step 1 In the main ASDM application window, choose **File** > **Reset Device to the Factory Default Configuration**.

The Reset Device to the Default Configuration dialog box appears.

Step 2 Enter the Management IP address of the management interface, instead of using the default address, 192.168.1.1. For an security appliance with a dedicated management interface, the interface is called "Management0/0." For other security appliance, the configured interface is Ethernet 1 and called "inside."

Step 3 Choose the Management (or Inside) Subnet Mask from the drop-down list.

Step 4 To save this configuration to internal flash memory, choose **File** > **Save Running Configuration to Flash**.

Selecting this option saves the running configuration to the default location for the startup configuration, even if you have previously configured a different location for Setting the System Time. When the configuration was cleared, this path was also cleared. The next time you reload the security appliance after restoring the factory configuration, the device boots from the first image in internal flash memory. If an image in internal flash memory does not exist, the security appliance does not boot.

Modes

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

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

Save the Running Configuration to a TFTP Server

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

To save the running configuration to a TFTP server, perform the following steps:

Step 1 In the main ASDM application window, choose **File** > **Save Running Configuration to TFTP Server**.

The Save Running Configuration to TFTP Server dialog box appears.

Step 2 Enter the TFTP server IP address and file path on the TFTP server in which the configuration file will be saved, and then click **Save Configuration**.

Note To configure default TFTP settings, choose **Configuration** > **Device Management** > **Management Access** > **File Access** > **TFTP Client**. After you have configured this setting, the TFTP server IP address and file path on the TFTP server appear automatically in this 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
•	•	•	•	•

Saving an Internal Log Buffer to Flash

This feature lets you save the internal log buffer to flash memory.

To save the internal log buffer to flash memory, perform the following steps:

Step 1 In the main ASDM application window, choose **File > Save Internal Log Buffer to Flash**.

The Enter Log File Name dialog box appears.

Step 2 Choose the first option to save the log buffer with the default filename, LOG-YYYY-MM-DD-hhmmss.txt.

Step 3 Choose the second option to specify a filename for the log buffer.

Step 4 Enter the filename for the log buffer, and then click **OK**.

Using the Command Line Interface

This feature provides a text-based tool for sending commands to the security appliance and viewing the results.

The commands you can enter with the CLI tool depend on your user privileges. See the "About Authorization" section on page 16-2 for more information. Review your privilege level in the status bar at the bottom of the main ASDM application window to ensure that you have the required privileges to execute privileged-level CLI commands.

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

To use the CLI tool, perform the following steps:

Step 1 In the main ASDM application window, choose **Tools > Command Line Interface**.

The Command Line Interface dialog box appears.

Step 2 Choose the type of command (single line or multiple line) that you want, and then choose the command from the drop-down list, or type it in the field provided.

Step 3 Click **Send** to execute the command.

Step 4 To enter a new command, click **Clear Response**, and then choose (or type) another command to execute.

Step 5 Check the **Enable context-sensitive help (?)** check box to provide context-sensitive help for this feature. Uncheck this check box to disable the context-sensitive help.

Step 6 After you have closed the Command Line Interface dialog box, if you changed the configuration, click **Refresh** to view the changes in ASDM.

Modes

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

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

Handling Command Errors

If an error occurs because you entered an incorrect command, the incorrect command is skipped and the remaining commands are processed. A message appears in the Response area to inform you whether or not any error occurred, as well as other related information.

Note ASDM supports almost all CLI commands. See the *Cisco ASA 5500 Series Command Reference* for a list of commands.

Using Interactive Commands

Interactive commands are not supported in the CLI tool. To use these commands in ASDM, use the **noconfirm** keyword if available, as shown in the following command:

```
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 CLI 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 most recent changes take effect.

To view other administrative sessions that are currently active on the same security appliance, choose **Monitoring > Properties > Device Access**.

Showing Commands Ignored by ASDM on the Device

This feature lets you show the list of commands that ASDM does not support. Typically, ASDM ignores them. ASDM does not change or remove these commands from your running configuration. See the "Unsupported Commands" section on page 2-3 for more information.

To display the list of unsupported commands for ASDM, perform the following steps:

Step 1 In the main ASDM application window, choose **Tools > Show Commands Ignored by ASDM on Device**.

Step 2 Click **OK** when you are done.

Modes

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

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

Using Diagnostic Tools

ASDM provides a set of diagnostic tools to help you in troubleshooting problems. This section includes the following topics:

- Tracing Packets with Packet Tracer, page 5-7
- Verifying ASA Configuration and Operation, and Testing Interfaces Using Ping, page 5-8
- Determining Packet Routing with Traceroute, page 5-11
- Sending an Administrator's Alert to Clientless SSL VPN Users, page 5-12
- Viewing and Copying Logged Entries with the ASDM Java Console, page 5-12
- Configuring and Running Captures with the Packet Capture Wizard, page 5-13

Tracing Packets with Packet Tracer

The packet tracer tool provides packet tracing for packet sniffing and network fault isolation, as well as detailed information about the packets and how they are processed by the security appliance. If a configuration command 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, the following message is displayed:

```
"packet dropped due to bad ip header (reason)."
```

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

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

To open the packet tracer, perform the following steps:

Step 1 In the main ASDM application window, choose **Tools > Packet Tracer**.

The Cisco ASDM Packet Tracer dialog box appears.

Step 2 Choose the source interface for the packet trace from the drop-down list.

Step 3 Specify the protocol type for the packet trace. Available protocol types include ICMP, IP, TCP, and UDP.

Step 4 Enter the source address for the packet trace in the Source IP Address field.

Step 5 Choose the source port for the packet trace from the drop-down list.

Step 6 Enter the destination IP address for the packet trace in the Destination IP Address field.

Step 7 Choose the destination port for the packet trace from the drop-down list.

Step 8 Click **Start** to trace the packet.

The Information Display Area shows detailed messages about the packet trace.

Note To display a graphical representation of the packet trace, check the **Show animation** check box.

Modes

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

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

Verifying ASA Configuration and Operation, and Testing Interfaces Using Ping

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

A ping is sent to an IP address and it returns a reply. This process enables network devices to discover, identify, and test each other.

The Ping tool uses ICMP (as described in RFC-777 and RFC-792) to define an echo request-and-reply transaction between two network devices. The echo request packet is sent to the IP address of a network device. The receiving device reverses the source and destination address and sends the packet back as the echo reply.

Administrators can use the ASDM Ping interactive diagnostic tool in these ways:

- Loopback testing of two interfaces—A ping may be initiated from one interface to another on the same security appliance, as an external loopback test to verify basic "up" status and operation of each interface.
- Pinging to an security appliance—The Ping tool can ping an interface on another security appliance 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 of functioning incorrectly. If the interface is configured correctly and an echo is not received, there may be problems with the device.
- Pinging to test intermediate communications—A ping may be initiated from an security appliance interface to a network device that is known to be functioning correctly and returning echo requests. If the echo is received, the correct operation of any intermediate devices and physical connectivity is confirmed.

Pinging from an Adaptive Security Appliance Interface

For basic testing of an interface, you can initiate a ping from a security appliance interface to a network device that you know is functioning correctly and returning replies via the intermediate communications path. For basic testing, make sure you do the following:

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

Pinging to an Adaptive Security Appliance Interface

When you try to ping to an security appliance interface, verify that the pinging response (ICMP echo reply*)* is enabled for that interface by choosing **Tools** > **Ping**. 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 Adaptive Security Appliance

To verify that other types of network traffic from "known good" sources is being passed through the security appliance, choose **Monitoring > Interfaces > Interface Graphs** or an SNMP management station.

To enable internal hosts to ping external hosts, configure ICMP access correctly for both the inside and outside interfaces by choosing **Configuration > Firewall > Objects > IP Names**.

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 necessarily because of no response from the IP address being pinged. Before using the Ping tool to ping from, to, or through an security appliance interface, perform the following basic checks:

- Verify that interfaces are configured by choosing **Configuration** > **Device Setup** > **Interfaces**.
- Verify that devices in the intermediate communications path, such as switches or routers, are correctly delivering other types of network traffic.
- Make sure that traffic of other types from "known good" sources is being passed by choosing **Monitoring** > **Interfaces** > **Interface Graphs**.

Using the Ping Tool

To use the Ping tool, perform the following steps:

Step 1 In the main ASDM application window, choose **Tools > Ping**.

The Ping dialog box appears.

Step 2 Enter the destination IP address for the ICMP echo request packets in the IP Address field.

Ping can also accept IPv6 addresses.

Note If a hostname has been assigned in the Configuration > Firewall > Objects > IP Names pane, you can use the hostname in place of the IP address.

Step 3 (Optional) Choose the security appliance interface that transmits the echo request packets from the drop-down list. If it is not specified, the security appliance checks the routing table to find the destination address and uses the required interface.

Step 4 Click **Ping** to send an ICMP echo request packet from the specified or default interface to the specified IP address and start the response timer.

The response appears in the Ping Output area. Three attempts are made to ping the IP address, and results display the following fields:

- The IP address of the device pinged or a device name, if available. The name of the device, if assigned Hosts/Networks, may be displayed, even if NO response is the result.
- When the ping is transmitted, a millisecond timer starts with a specified maximum, or timeout value. This timer is useful for testing the relative response times of different routes or activity levels.
- Example Ping output:

```
Sending 5, 100-byte ICMP Echos to out-pc, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
If the ping fails, the output is as follows:
Sending 5, 100-byte ICMP Echos to 10.132.80.101, timeout is 2 seconds:
?????
Success rate is 0 percent (0/5)
```

Step 5 To enter a new IP address, click **Clear Screen** to remove the previous response from the Ping output area.

Modes

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

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

Determining Packet Routing with Traceroute

The Traceroute tool helps you to determine the route that packets will take to their destination. The tool prints the result of each probe sent. Every line of output corresponds to a TTL value in increasing order. The following table lists the output symbols printed by this tool.

Output Symbol	Description
*	No response was received for the probe within the timeout period.
nn msec	For each node, the round-trip time (in milliseconds) for the specified number of probes.
!N.	ICMP network unreachable.
!H	ICMP host unreachable.
!P	ICMP unreachable.
!A	ICMP administratively prohibited.
?	Unknown ICMP error.

To use the Traceroute tool, perform the following steps:

Step 1 In the main ASDM application window, choose **Tools > Traceroute**.

The Traceroute dialog box appears.

Step 2 Enter the name of the host to which the route is traced. If the hostname is specified, define it by choosing **Configuration > Firewall > Objects > IP Names**, or configure a DNS server to enable this tool to resolve the hostname to an IP address.

Step 3 Enter the amount of time in seconds to wait for a response before the connection times out. The default is three seconds.

Step 4 Type the destination port used by the UDP probe messages. The default is 33434.

Step 5 Enter the number of probes to be sent at each TTL level. The default is three.

Step 6 Specify the minimum and maximum TTL values for the first probes. The minimum default is one, but it can be set to a higher value to suppress the display of known hops. The maximum default is 30. The traceroute terminates when the packet reaches the destination or when the maximum value is reached.

Step 7 Check the **Specify source interface or IP address** check box. Choose the source interface or IP address for the packet trace from the drop-down list. This IP address must be the IP address of one of the interfaces. In transparent mode, it must be the management IP address of the security appliance.

Step 8 Check the **Reverse Resolve** check box to have the output display the names of hops encountered if name resolution is configured. Leave this check box unchecked to have the output display IP addresses.

Step 9 Check the **Use ICMP** check box to specify the use of ICMP probe packets instead of UDP probe packets.

Step 10 Click **Trace Route** to start the traceroute.

The Traceroute Output area displays detailed messages about the traceroute results.

Step 11 Click **Clear Output** to start a new traceroute.

Modes

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

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

Sending an Administrator's Alert to Clientless SSL VPN Users

This feature lets you send an alert message to clientless SSL VPN users (for example, about connection status).

To send an alert message, perform the following steps:

Step 1 In the main ASDM application window, choose **Tools > Administrator's Alert Message to Clientless SSL VPN Users**.

The Administrator's Alert Message to Clientless SSL VPN Users dialog box appears.

Step 2 Enter the new or edited alert content that you want to send, and then click **Post Alert**.

Step 3 To remove current alert content and enter new alert content, click **Cancel Alert**.

Modes

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

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

Viewing and Copying Logged Entries with the ASDM Java Console

You can use the ASDM Java console to view and copy logged entries in a text format, which can help you troubleshoot ASDM errors.

To access the ASDM Java Console, perform the following steps:

Step 1 In the main ASDM application window, choose **Tools > ASDM Java Console**.

Step 2 To show the virtual machine memory statistics, enter **m** in the console.

Step 3 To perform garbage collection, enter **g** in the console.

Step 4 To monitor memory usage, open the Windows Task Manager and double-click the **asdm_launcher.exe** file.

Note The maximum memory allocation allowed is 256 MB.

Modes

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

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

Configuring and Running Captures with the Packet Capture Wizard

You can use the Packet Capture Wizard to configure and run captures for troubleshooting errors. The captures can use access lists to limit the type of traffic captured, the source and destination addresses and ports, and one or more interfaces. The wizard runs one capture on each of the ingress and egress interfaces. You can save the captures on your PC to examine them in a packet analyzer.

Note This tool does not support clientless SSL VPN capture.

To configure and run captures, perform the following steps:

Step 1 In the main ASDM application window, choose **Wizards > Packet Capture Wizard**.

The Overview of Packet Capture screen appears, with a list of the tasks through which the wizard will guide you to complete.

Step 2 Click **Next** to display the Ingress Traffic Selector screen.

Step 3 Choose the ingress interface from the drop-down list.

Step 4 In the Packet Match Criteria area, do one of the following:

- To specify the access list to use for matching packets, click the **Specify access-list** radio button, and then choose the access list from the Select access list drop-down list. To add a previously configured access list to the current drop-down list, click **Manage** to display the ACL Manager pane. Choose an access list, and click **OK**.
- To specify packets parameters, click the **Specify Packet Parameters** radio button.

Step 5 Click **Next** to display the Ingress Traffic Selector screen. For more information, see the "Ingress Traffic Selector" section on page 5-15.

Step 6 Enter the source host IP address and choose the network IP address from the drop-down list.

Step 7 Enter the destination host IP address and choose the network IP address from the drop-down list.

Step 8 Choose the protocol type to capture from the drop-down list. Available protocol types to capture are ah, eigrp, esp, gre, icmp, icmp6, igmp, igrp, ip, ipinip, nos, ospf, pcp, pim, snp, tcp, or udp.

Step 9 Click **Next** to display the Egress Traffic Selector screen. For more information, see the "Egress Traffic Selector" section on page 5-16.

Step 10 Choose the egress interface from the drop-down list.

Step 11 Enter the source host IP address and choose the network IP address from the drop-down list.

Step 12 Enter the destination host IP address and choose the network IP address from the drop-down list.

Note The source port services, destination port services, and ICMP type are read-only based on the choices you made in the Ingress Traffic Selector screen.

Step 13 Click **Next** to display the Buffers & Captures screen. For more information, see the "Buffers" section on page 5-16.

Step 14 In the Capture Parameters area, to obtain the latest capture every 10 seconds automatically, check the **Get capture every 10 seconds** check box. By default, this capture uses the circular buffer.

Step 15 In the Buffer Parameters area, you specify the buffer size and packet size. The buffer size is the maximum amount of memory that the capture can use to store packets. The packet size is the longest packet that the capture can hold. We recommend that you use the longest packet size to capture as much information as possible.

a. Enter the packet size. The valid size ranges from 14 - 1522 bytes.

b. Enter the buffer size. The valid size ranges from 1534 - 33554432 bytes.

c. Check the **Use circular buffer** check box to store captured packets.

Note When you choose this setting, if all the buffer storage is used, the capture starts overwriting the oldest packets.

Step 16 Click **Next** to display the Summary screen, which shows the traffic selectors and buffer parameters that you have entered. For more information, see the "Summary" section on page 5-17.

Step 17 Click **Next** to display the Run Captures screen, and then click **Start** to begin capturing packets. Click **Stop** to end the capture. For more information, see the "Run Captures" section on page 5-17.

Step 18 Click **Get Capture Buffer** to determine how much buffer space you have remaining. Click **Clear Buffer on Device** to remove the current content and allow room in the buffer to capture more packets.

Step 19 Click **Save captures** to display the Save Capture dialog box. Choose the format in which you want to include the captures: **ASCII** or **PCAP**. You have the option of saving either the ingress capture, the egress capture, or both.

Step 20 To save the ingress packet capture, click **Save Ingress Capture** to display the Save capture file dialog box. Specify the storage location on your PC, and click **Save**.

Step 21 Click **Launch Network Sniffer Application** to start the packet analysis application specified in Tools > Preferences for analyzing the ingress capture.

Step 22 To save the egress packet capture, click **Save Egress Capture** to display the Save capture file dialog box. Specify the storage location on your PC, and click **Save**.

Step 23 Click **Launch Network Sniffer Application** to start the packet analysis application specified in Tools > Preferences for analyzing the egress capture.

Step 24 Click **Close,** and then click **Finish** to exit the wizard.

Modes

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

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

Ingress Traffic Selector

To configure the ingress interface, source and destination hosts/networks, and the protocol for packet capture, perform the following steps:

Step 1 Enter the ingress interface name.

Step 2 Enter the ingress source host and network.

Step 3 Enter the ingress destination host and network.

Step 4 Enter the protocol type to capture. Available protocols include ah, eigrp, esp, gre, icmp, icmp6, igmp, igrp, ip, ipinip, nos, ospf, pcp, pim, snp, tcp, or udp.

a. Enter the ICMP type for the ICMP protocol only. Available types include all, alternate address, conversion-error, echo, echo-reply, information-reply, information-request, mask-reply, mask-request, mobile-redirect, parameter-problem, redirect, router-advertisement, router-solicitation, source-quench, time-exceeded, timestamp-reply, timestamp-request, traceroute, or unreachable.

b. Specify the source and destination port services for the TCP and UDP protocols only. Available options include the following:

   - To include all services, choose All Services.
   - To include a service group, choose Service Groups.
   - To include a specific service, choose one of the following: aol, bgp, chargen, cifx, citrix-ica, ctiqbe, daytime, discard, domain, echo, exec, finger, ftp, ftp-data, gopher, h323, hostname, http, https, ident, imap4, irc, kerberos, klogin, kshell, ldap, ldaps, login, lotusnotes, lpd, netbios-ssn, nntp, pcanywhere-data, pim-auto-rp, pop2, pop3, pptp, rsh, rtsp, sip, smtp, sqlnet, ssh, sunrpc, tacacs, talk, telnet, uucp, or whois.

Modes

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

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

Egress Traffic Selector

To configure the egress interface, source and destination hosts/networks, and source and destination port services for packet capture, perform the following steps:

Step 1 Enter the egress interface name.

Step 2 Enter the egress source host and network.

Step 3 Enter the egress destination host and network.

The protocol type selected during the ingress configuration is already listed.

Modes

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

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

Buffers

To configure the packet size, buffer size, and use of the circular buffer for packet capture, perform the following steps.

Step 1 Enter the longest packet that the capture can hold. Use the longest size available to capture as much information as possible.

Step 2 Enter the maximum amount of memory that the capture can use to store packets.

Step 3 Use the circular buffer to store packets. When the circular buffer has used all of the buffer storage, the capture will overwrite the oldest packets first.

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 shows the traffic selectors and the buffer parameters for the packet capture selected in the previous wizard screens.

Modes

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

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

Run Captures

To start and stop the capture session, view the capture buffer, launch a network analyzer application, save packet captures, and clear the buffer, perform the following steps:

Step 1 To begin the packet capture session on a selected interface, click **Start**.

Step 2 To stop the packet capture session on a selected interface, click **Stop**.

Step 3 To obtain a snapshot of the captured packets on the interface, click **Get Capture Buffer**.

Step 4 To show the capture buffer on the ingress interface, click **Ingress**.

Step 5 To show the capture buffer on the egress interface, click **Egress**.

Step 6 To clear the buffer on the device, click **Clear Buffer on Device**.

Step 7 To start the packet analysis application for analyzing the ingress capture or the egress capture specified in Tools > Preferences, click **Launch Network Sniffer Application**.

Step 8 To save the ingress and egress captures in either ASCII or PCAP format, click **Save Captures**.

Modes

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

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

Save Captures

To save the ingress and egress packet captures to ASCII or PCAP file format for further packet analysis, perform the following steps:

Step 1 To save the capture buffer in ASCII format, click **ASCII**.

Step 2 To save the capture buffer in PCAP format, click **PCAP**.

Step 3 To specify a file in which to save the ingress packet capture, click **Save ingress capture**.

Step 4 To specify a file in which to save the egress packet capture, click **Save egress capture**.

Modes

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

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

Using File Management Tools

ASDM provides a set of file management tools to help you perform basic file management tasks. This section includes the following topics:

- Managing Files, page 5-19
- Managing Mount Points, page 5-20
- Adding or Editing a CIFS/FTP Mount Point, page 5-20
- Accessing a CIFS Mount Point, page 5-21
- Upgrading Software from Your Local Computer, page 5-22
- Transferring Files, page 5-23
- Upgrading Software from the Cisco.com Wizard, page 5-24
- Scheduling a System Restart, page 5-26
- Backing Up and Restoring Configurations, Images, and Profiles, page 5-27

Managing Files

The File Management tool lets you view, move, copy, and delete files stored in flash memory, transfer files, and to manage files on remote storage devices (mount points).

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

To use the file management tools, perform the following steps:

Step 1 In the main ASDM application window, choose **Tools > File Management**.

The File Management dialog box appears.

- The Folders pane displays the available folders on disk.
- Flash Space shows the total amount of flash memory and how much memory is available.
- The Files area displays the following information about files in the selected folder:
  - Path
  - Filename
  - Size (bytes)
  - Time Modified
  - Status, which indicates whether a selected file is designated as a boot configuration file, boot image file, ASDM image file, SVC image file, CSD image file, or APCF image file.

Step 2 Click **View** to display the selected file in your browser.

Step 3 Click **Cut** to cut the selected file for pasting to another directory.

Step 4 Click **Copy** to copy the selected file for pasting to another directory.

Step 5 Click **Paste** to paste the copied file to the selected destination.

Step 6 Click **Delete** to remove the selected file from flash memory.

Step 7 Click **Rename** to rename a file.

Step 8 Click **New Directory** to create a new directory for storing files.

Step 9 Click **File Transfer** to open the File Transfer dialog box. See the "Transferring Files" section on page 5-23 for more information.

Step 10 Click **Mount Points** to open the Manage Mount Points dialog box. See the "Managing Mount Points" section on page 5-20 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
•	•	•	—	•

Managing Mount Points

This feature lets you configure remote storage (mount points) for network file systems using a CIFS or FTP connection. The dialog box lists the mount-point name, connection type, server name or IP address, and the enabled setting (yes or no). You can add, edit, or delete mount points. See the "Adding or Editing a CIFS/FTP Mount Point" section on page 5-20 for more information. You can access a CIFS mount point after it has been created. For more information, see Accessing a CIFS Mount Point, page 5-21.

Note On a PIX 535 security appliance in single, routed mode, the Managing Mount Point feature is not available.

Adding or Editing a CIFS/FTP Mount Point

To add a CIFS mount point, perform the following steps:

Step 1 Click **Add**, and then choose **CIFS Mount Point**.

The Add CIFS Mount Point dialog box appears.

The Enable mount point check box is automatically checked, which is the default setting.

Step 2 Enter the mount-point name, server name or IP address, and share name in the applicable fields.

Step 3 In the Authentication section, enter the NT domain, username and password, and then confirm the password.

Step 4 Click **OK**.

To add an FTP mount point, perform the following steps:

Step 1 Click **Add**, and then choose **FTP Mount Point**.

The Add FTP Mount Point dialog box appears.

The Enable mount point check box is automatically checked, which is the default setting.

Step 2 Enter the mount-point name and the server name or IP address in the applicable fields.

Step 3 In the FTP Mount Options area, click the **Active Mode** or **Passive Mode** option.

Step 4 Enter the path to mount the remote storage.

Step 5 In the Authentication area, enter the NT domain, username and password, and then confirm the password.

Step 6 Click **OK**.

To edit a CIFS mount point, perform the following steps:

Step 1 Choose the CIFS mount-point you want to modify, and click **Edit**.

The Edit CIFS Mount Point dialog box appears.

Note You cannot change the CIFS mount-point name.

Step 2 Make the changes to the remaining settings, and click **OK** when you are done.

To edit an FTP mount point, perform the following steps:

Step 1 Choose the FTP mount-point you want to modify, and click **Edit**.

The Edit FTP Mount Point dialog box appears.

Note You cannot change the FTP mount-point name.

Step 2 Make the changes to the remaining settings, and click **OK** when you are done.

Accessing a CIFS Mount Point

To access a CIFS mount point after it has been created, perform the following steps:

Step 1 Start the security appliance CLI.

Step 2 Create the mount by entering the **mount <name of mount> type cifs** command.

Step 3 Enter the **show run mount** command.

The following output appears:

Note In this example, win2003 is the name of the mount.

```
server kmmwin2003
share sharefolder
username webvpnuser2
password ********
status enable
```

Step 4 Enter the **dir** command to list all enabled mounts as subdirectories, which is similar to mounting a drive on the Windows PC. For example, in the following output, FTP2003: , FTPLINUX:, and win2K: are configured mounts.

The following is sample output from the **dir** command:

```
FTP2003: Directory or file name
FTPLINUX: Directory or file name
WIN2003: Directory or file name
all-filesystems List files on all filesystems
disk0: Directory or file name
disk1: Directory or file name
flash: Directory or file name
system: Directory or file name
win2K: Directory or file name
```

Step 5 Enter the **dir** command for that mount (for example, **dir WIN2003**), and copy files to and from flash (disk0:) to any of the listed mounts.

The following is sample output from the **dir WIN2003** command.

```
Directory of WIN2003:/
---- 14920928 08:33:36 Apr 03 2009 1_5_0_01-windows-i586-p.exe
---- 33 11:27:16 Jun 07 2007 AArenameIE70
---- 28213021 15:15:22 Apr 03 2009 atest2(3).bin
---- 61946730 12:09:40 Mar 17 2009 atest2.bin
---- 5398366 14:52:10 Jul 28 2008 atest222.bin
---- 2587728 10:07:44 Dec 06 2005 cCITRIXICA32t.exe
---- 1499578 15:26:50 Dec 02 2005 ccore.exe
---- 61946728 11:40:36 Dec 09 2005 CIFSTESTT.bin
---- 2828 13:46:04 May 11 2009 ClientCert.pfx
d--- 16384 14:48:28 Mar 20 2007 cookiefolder
---- 4399 15:58:46 Jan 06 2006 Cookies.plist
---- 2781710 12:35:00 Dec 12 2006 coreftplite1.3.exe
---- 0 10:22:52 Jul 13 2007 coreftplite1.3.exe.download
---- 245760 15:13:38 Dec 21 2005 Dbgview.exe
---- 1408249 11:01:34 Dec 08 2005 expect-5.21r1b1-setup.exe
d--- 16384 14:49:14 Jul 28 2008 folder157
---- 101 09:33:48 Dec 12 2005 FxSasser.log
---- 2307104 09:54:12 Dec 12 2005 ica32t.exe
---- 8732552 10:14:32 Apr 29 2009 iclientSetup_IFen_flex51.exe
d--- 16384 08:32:46 Apr 03 2009 IE8withVistaTitan
---- 15955208 08:34:18 Aug 14 2007 j2re.exe
---- 16781620 13:38:22 Jul 23 2008 jre-1_5_0_06-windows-i586-p.exe
<--- More --->
```

Upgrading Software from Your Local Computer

The Upgrade Software from Local Computer tool lets you upload an image file from your PC to the flash file system to upgrade the security appliance.

To upgrade software from your PC, perform the following steps:

Step 1 In the main ASDM application window, choose **Tools > Upgrade Software from Local Computer**.

The Upgrade Software from Local Computer dialog box appears.

Step 2 Choose the image file to upload from the drop-down list.

Step 3 Enter the local path to the file on your PC or click **Browse Local Files** to find the file on your PC.

Step 4 Enter the path to the flash file system or click **Browse Flash** to find the directory or file in the flash file system.

Step 5 Click **Image to Upload**. The uploading process might take a few minutes; make sure you wait until it is finished.

Modes

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

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

Transferring Files

The File Transfer tool lets you transfer files from either a local or remote location. You can transfer a local file on your computer or a flash file system to and from the security appliance. You can transfer a remote file to and from the security appliance using HTTP, HTTPS, TFTP, FTP, or SMB.

To transfer files between your local computer and a flash file system, perform the following steps:

Step 1 In the main ASDM application window, choose **Tools > File Management**.

The File Management dialog box appears.

Step 2 Click the down arrow next to **File Transfer**, and then click **Between Local PC and Flash**.

The File Transfer dialog box appears.

Step 3 Select and *drag* the file(s) from either your local computer or the flash file system that you want to upload or download to the desired location. Alternatively, select the file(s) from either your local computer or the flash file system that you want to upload or download, and click the right arrow or left arrow to transfer the file(s) to the desired location.

Step 4 Click **Close** when you are done.

To transfer files between a remote server and a flash file system, perform the following steps:

Step 1 In the main ASDM application window, choose **Tools > File Management**.

The File Management dialog box appears.

Step 2 Click the down arrow from the File Transfer drop-down list, and then click **Between Remote Server and Flash**.

The File Transfer dialog box appears.

Step 3 To transfer a file from a remote server, click the **Remote server** option.

Step 4 Define the source file to be transferred.

a. Choose the path to the location of the file, including the IP address of the server.

Note File transfer supports IPv4 and IPv6 addresses.

b. Enter the type (if the path is FTP) or the port number (if the path is HTTP or HTTPS) of the remote server. Valid FTP types are the following:

- 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

Step 5 To transfer the file from the flash file system, click the **Flash file system** option.

Step 6 Enter the path to the location of the file or click **Browse Flash** to find the file location.

Step 7 In addition, you can copy a file from your startup configuration, running configuration, or an SMB file system through the CLI. For instructions about using the **copy** command, see the *Cisco ASA 5500 Series Configuration Guide using the CLI*.

Step 8 Define the destination of the file to be transferred.

a. To transfer the file to the flash file system, choose the **Flash file system** option.

b. Enter the path to the location of the file or click **Browse Flash** to find the file location.

Step 9 To transfer a file to a remote server, choose the **Remote server** option.

a. Enter the path to the location of the file.

b. For FTP transfers, enter the type. Valid types are the following:

- 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

Step 10 Click **Transfer** to start the file transfer.

The Enter Username and Password dialog box appears.

Step 11 Enter the username, password, and domain (if required) for the remote server.

Step 12 Click **OK** to continue the file transfer.

The file transfer process might take a few minutes; make sure that you wait until it is finished.

Step 13 Click **Close** when the file transfer is finished.

Modes

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

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

Upgrading Software from the Cisco.com Wizard

The Upgrade Software from Cisco.com Wizard lets you automatically upgrade the ASDM and security appliance to more current versions.

Note This feature is not available in the user or admin context mode in a single security context.

In this wizard, you can do the following:

- Download the list of available releases from Cisco.com.
- Choose an ASDM image file or ASA image file to upgrade.
- Upgrade the selected images.
- Restart the security appliance if you have upgraded the ASA image (optional).

Note You must upgrade incrementally from one version to the next (for example, from Version 6.1 to 6.2, from Version 6.1(3) to 6.1(5), and so on). You cannot upgrade from Version 5.2(3) to 6.1(3).

To upgrade software from Cisco.com, perform the following steps:

Step 1 In the main ASDM application window, choose **Tools > Upgrade Software from Cisco.com**.

The Upgrade Software from Cisco.com Wizard appears. The Overview screen describes the steps in the image upgrade process.

Step 2 Click **Next** to display the Authentication screen.

Step 3 Enter your assigned Cisco.com user name and the Cisco.com password, and then click **Next** to display the Image Selection screen.

Step 4 Choose one or both of the two options listed.

- Check the **Upgrade the ASA version** check box to specify the most current security appliance image to which you want to upgrade.
- Check the **Upgrade the ASDM version** check box to specify the most current ASDM version to which you want to upgrade.

Note If the ASA version list or the ASDM version list is empty, a statement appears informing you that no new ASA or ASDM images are available. If you see this statement, you can exit the wizard.

Step 5 Click **Next** to display the Selected Images screen.

Step 6 Verify that the image file you have selected is the correct one, and then click **Next** to start the upgrade.

The wizard indicates that the upgrade will take a few minutes. You can then view the status of the upgrade as it progresses.

The Results screen appears, which provides additional details, such as the upgrade status (success or failure) or a request asking if you want to save the configuration and restart the security appliance.

If the upgrade succeeded, an option to save the configuration and restart the security appliance appears.

Step 7 Click **Yes**.

For the upgrade versions to take effect, you must save the configuration, restart the security appliance, and restart ASDM.

Note You do not need to restart the wizard after you have completed one incremental upgrade. You can return to Step 3 of the wizard to upgrade to the next higher version, if any.

Step 8 Click **Finish** to exit the wizard when the upgrade is finished.

Modes

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

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

Scheduling a System Restart

The System Reload tool lets you schedule a system restart or cancel a pending restart.

To schedule a system restart, perform the following steps:

Step 1 In the main ASDM application window, choose **Tools > System Reload**.

Step 2 In the Reload Scheduling area, define the following settings:

a. For the Configuration State, choose either to save or discard the running configuration at restart time.

b. For the Reload Start Time, choose from the following options:

- Click **Now** to perform an immediate restart.
- Click **Delay by** to delay the restart by a specified amount of time. Enter the time before the restart begins in hours and minutes or only minutes.
- Click **Schedule at** to schedule the restart to occur at a specific time and date. Enter the time of day the restart is to occur, and select the date of the scheduled restart.

c. In the Reload Message field, enter a message to send to open instances of ASDM at restart time.

d. Check the **On reload failure force immediate reload after** check box to show the amount of time elapsed in hours and minutes or only minutes before a restart is attempted again.

e. Click **Schedule Reload** to schedule the restart as configured.

The Reload Status area displays the status of the restart.

Step 3 Choose one of the following:

- Click **Cancel Reload** to stop a scheduled restart.
- Click **Refresh** to refresh the Reload Status display after a scheduled restart is finished.
- Click **Details** to display the results of a scheduled restart.

Modes

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

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

Backing Up and Restoring Configurations, Images, and Profiles

The Backup and Restore features options on the Tools menu let you back up and restore the security appliance configuration, Cisco Secure Desktop image, and SSL VPN Client images and profiles.

ASDM lets you choose the file types to back up, compresses them into a single zip file, then transfers the zip file to the directory you choose on your computer. Similarly, to restore files, you choose the source zip file on your computer and then choose the file types to be restored.

This section includes the following topics:

- Backing Up Configurations, page 5-27
- Restoring Configurations, page 5-28

Backing Up Configurations

To back up configurations and images to a .zip file to be transferred to your local computer, perform the following steps:

Step 1 Create a folder on your computer to store backup files so they will be easy to find if you have to restore them later.

Step 2 Choose **Tools** > **Backup Configurations**.

The Backup Configurations dialog box appears.

By default, all files are checked and will be backed up if they are available. If you want to back up all of the files in the list, go to Step 5.

Step 3 Uncheck the **Backup All** check box if you want to specify the configurations to back up.

Step 4 Check the options to customize the backup.

Step 5 Click **Browse Local Directory**.

The Select dialog box appears.

Step 6 Choose the path on your computer to specify the target destination for the zip file to package the backup.

Step 7 Click **Select**.

The path appears in the Local File field.

Step 8 Enter the name of the destination backup file after the path.

Step 9 Click **Backup**.

A status window appears. When the backup completes, the status window closes and the Backup Statistics dialog box appears.

This dialog box shows the status of each backup.

Note Backup "failure messages" are most likely the consequence of no configuration present for the types indicated.

Step 10 Click **OK** to close the Backup Statistics 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
•	•	•	—	•

Restoring Configurations

You can specify configurations and images to restore from a zip file on your local computer. Before proceeding, note the following restrictions:

- The zip file you restore must be created using the Tools > Backup Configurations option.
- The DAP configuration may depend on a specific running configuration, URL list, and CSD configuration.
- The CSD configuration may depend on the version of the CSD image.
- You can restore components, images, and configurations using backups made from the same adaptive security appliance type. You must start with a basic configuration that allows ASDM access. Choose **Tools** > **Restore Configurations** to restore the applicable items, then use the CLI to restore the backup running configuration, as follows: Unzip the running-config.cfg file created during a previous backup, and transfer it to the adaptive security appliance file system. On the CLI, enter **copy running-config.cfg running-config**. To preserve the configuration, enter **copy running-config start-config**.

To restore selected elements of the security appliance configuration, Cisco Secure Desktop image, or SSL VPN Client images and profiles, perform the following steps:

Step 1 Choose **Tools > Restore Configurations**.

The first Restore Configurations dialog box appears.

Note Later in the procedure, you have an opportunity to choose the configuration elements to restore; this window lets you choose the file from which to restore them.

Step 2 Click **Browse Local Directory**, choose the zip file on your local computer that includes the configuration to restore, then click **Select**.

The path and the zip filename appear in the Local File field.

Step 3 Click **Next**.

The second Restore Configuration dialog box appears.

By default, all files are checked; ASDM restores them if they are available.

Step 4 Use the default options, or uncheck these check boxes and check the check boxes for specific configurations and images that you want to restore.

Step 5 Click **Restore**.

ASDM displays a status window until the restore operation completes.

Modes

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

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

CHAPTER 6

Configuring the Transparent or Routed Firewall

This chapter describes how to configure the firewall mode, routed or transparent, and how to customize transparent firewall operation.

Note In multiple context mode, you cannot set the firewall mode separately for each context; you can only set the firewall mode for the entire security appliance.

This chapter includes the following sections:

Configuring the Firewall Mode

This section describes routed and transparent firewall mode, and how to set the mode. This section includes the following topics:

Information About the Firewall Mode

This section describes routed and transparent firewall mode, and includes the following topics:

Information About Routed Firewall Mode

In routed mode, the security appliance is considered to be a router hop in the network. It can use OSPF or RIP (in single context mode). Routed mode supports many interfaces. Each interface is on a different subnet. You can share interfaces between contexts.

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. 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.

Information About Transparent Firewall Mode

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.

This section describes transparent firewall mode, and includes the following topics:

- Transparent Firewall Network, page 6-2
- Allowing Layer 3 Traffic, page 6-2
- Allowed MAC Addresses, page 6-2
- Passing Traffic Not Allowed in Routed Mode, page 6-3
- MAC Address vs. Route Lookups, page 6-3
- Using the Transparent Firewall in Your Network, page 6-4

Transparent Firewall Network

The security appliance connects the same network on its inside and outside interfaces. Because the firewall is not a routed hop, you can easily introduce a transparent firewall into an existing network.

Allowing Layer 3 Traffic

IPv4 and IPv6 traffic is allowed through the transparent firewall automatically from a higher security interface to a lower security interface, without an access list. ARPs are allowed through the transparent firewall in both directions without an access list. ARP traffic can be controlled by ARP inspection. For Layer 3 traffic travelling from a low to a high security interface, an extended access list is required on the low security interface. See Chapter 21, “Configuring Access Rules and ACLs,” for more information.

Allowed MAC Addresses

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

Passing Traffic Not Allowed in Routed Mode

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 any traffic through using either an extended access list (for IP traffic) or an EtherType access list (for non-IP traffic).

Note The transparent mode security appliance does not pass CDP packets packets, or any packets that do not have a valid EtherType greater than or equal to 0x600. For example, you cannot pass IS-IS packets. An exception is made for BPDUs, which are supported.

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.

MAC Address vs. Route Lookups

When the security appliance runs in transparent mode without NAT, 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.

An exception to this rule is when you use voice inspections and the endpoint is at least one hop away from the security appliance. For example, if you use the transparent firewall between a CCM and an H.323 gateway, and there is a router between the transparent firewall and the H.323 gateway, then you need to add a static route on the security appliance for the H.323 gateway for successful call completion.

If you use NAT, then the security appliance uses a route lookup instead of a MAC address lookup. In some cases, you will need static routes. For example, if the real destination address is not directly-connected to the security appliance, then you need to add a static route on the security appliance for the real destination address that points to the downstream router.

Using the Transparent Firewall in Your Network

Figure 6-1 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 6-1 Transparent Firewall Network

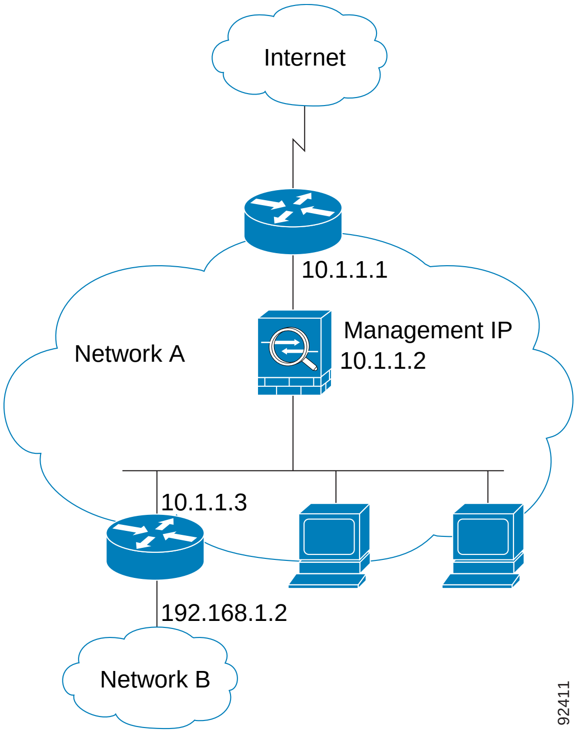

Licensing Requirements for the Firewall Mode

The following table shows the licensing requirements for this feature.

Model	License Requirement
All models	Base License.

Default Settings

The default mode is routed mode.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

- The firewall mode is set for the entire system and all contexts; you cannot set the mode individually for each context.
- For multiple context mode, set the mode in the system execution space.
- When you change modes, the security appliance clears the running configuration because many commands are not supported for both modes. 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 might 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.

Transparent Firewall Guidelines

Follow these guidelines when planning your transparent firewall network:

- For IPv4, a management IP address is required for both management traffic and for traffic to pass through the security appliance. 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).

  For IPv6, at a minimum you need to configure link-local addresses for each interface for through traffic. For full functionality, including the ability to manage the security appliance, you need to configure a global IP address for the device.

  You can configure an IP address (both IPv4 and IPv6) for the Management 0/0 or Management 0/1 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.

IPv6 Guidelines

Supports IPv6.

Additional Guidelines and Limitations

- When you change modes, the security appliance clears the running configuration because many commands are not supported for both modes. The startup configuration remains unchanged. If you reload without saving, then the startup configuration is loaded, and the mode reverts back to the original setting. See the "Setting the Firewall Mode" section on page 6-7 for information about backing up your configuration file.
- 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 appears later in the configuration, the security appliance clears all the preceding lines in the configuration.

Unsupported Features in Transparent Mode

Table 6-1 lists the features are not supported in transparent mode.

Table 6-1 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.
Multicast IP routing	You can allow multicast traffic through the security appliance by allowing it in an extended access list.
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. SSL VPN is also not supported.

Setting the Firewall Mode

This section describes how to change the firewall mode using the CLI. You cannot change the mode in ASDM.

Note We recommend that you set the firewall mode before you perform any other configuration because changing the firewall mode clears the running configuration.

Prerequisites

When you change modes, the security appliance clears the running configuration (see the "Guidelines and Limitations" section on page 6-5 for more information).

- 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.
- Use the CLI at the console port to change the mode. If you use any other type of session, including the ASDM Command Line Interface tool or SSH, you will be disconnected when the configuration is cleared, and you will have to reconnect to the security appliance using the console port in any case.

Detailed Steps

Command	Purpose

`firewall transparent`
Example:
`hostname(config)# firewall transparent`
	Sets the firewall mode to transparent. Enter this command in the system execution space for multiple context mode. To change the mode to routed, enter the **no firewall transparent** command.
	This command also appears in each context configuration for informational purposes only; you cannot enter this command in a context.
	Note You are not prompted to confirm the firewall mode change; the change occurs immediately.

Feature History for Firewall Mode

Table 6-2 lists the release history for this feature.

Table 6-2 Feature History for Firewall Mode

Feature Name	Releases	Feature Information
Transparent firewall mode	7.0(1)	A transparent firewall 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 following commands were introduced: **firewall transparent**, **show firewall**.

Configuring ARP Inspection for the Transparent Firewall

This section describes ARP inspection and how to enable it, and includes the following topics:

Information About 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.

Licensing Requirements for ARP Inspection

The following table shows the licensing requirements for this feature.

Model	License Requirement
All models	Base License.

Default Settings

By default, all ARP packets are allowed through the security appliance.

If you enable ARP inspection, the default setting is to flood non-matching packets.

Guidelines and Limitations

Context Mode Guidelines

- Supported in single and multiple context mode.
- In multiple context mode, configure ARP inspection within each context.

Firewall Mode Guidelines

Supported only in transparent firewall mode. Routed mode is not supported.

Configuring ARP Inspection

This section describes how to configure ARP inspection, and includes the following topics:

- Task Flow for Configuring ARP Inspection, page 6-9
- Adding a Static ARP Entry, page 6-9
- Enabling ARP Inspection, page 6-10

Task Flow for Configuring ARP Inspection

Follow these steps to configure ARP Inspection:

Step 1 Add static ARP entries according to the "Adding a Static ARP Entry" section on page 6-9. ARP inspection compares ARP packets with static ARP entries in the ARP table, so static ARP entries are required for this feature.

Step 2 Enable ARP inspection according to the "Enabling ARP Inspection" section on page 6-10.

Adding a Static ARP Entry

ARP inspection compares ARP packets with static ARP entries in the ARP 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.

Detailed Steps

Step 1 Choose the **Configuration** > **Device Setup** > **ARP** > **ARP Static Table** pane.

Step 2 (Optional) To set the ARP timeout for *dynamic* ARP entries, enter a value in the ARP Timeout field.

This field 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.

Step 3 Click **Add**.

The Add ARP Static Configuration dialog box appears.

Step 4 From the Interface drop-down list, choose the interface attached to the host network.

Step 5 In the IP Address field, enter the IP address of the host.

Step 6 In the MAC Address field, enter the MAC address of the host; for example, 00e0.1e4e.3d8b.

Step 7 To perform proxy ARP for this address, check the **Proxy ARP** check box.

If the security appliance receives an ARP request for the specified IP address, then it responds with the specified MAC address.

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

What to Do Next

Enable ARP inspection according to the "Enabling ARP Inspection" section on page 6-10.

Enabling ARP Inspection

This section describes how to enable ARP inspection.

Detailed Steps

Step 1 Choose the **Configuration** > **Device Setup** > **ARP** > **ARP Inspection** pane.

Step 2 Choose the interface row on which you want to enable ARP inspection, and click **Edit**.

The Edit ARP Inspection dialog box appears.

Step 3 To enable ARP inspection, check the **Enable ARP Inspection** check box.

Step 4 (Optional) To flood non-matching ARP packets, check the **Flood ARP Packets** check box.

By default, 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 uncheck this check box, all non-matching packets are dropped, which restricts ARP through the security appliance to only static entries.

Note The Management 0/0 or 0/1 interface or subinterface, if present, never floods packets even if this parameter is set to flood.

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

Feature History for ARP Inspection

Table 6-2 lists the release history for this feature.

Table 6-3 Feature History for ARP Inspection

Feature Name	Releases	Feature Information
ARP inspection	7.0(1)	ARP inspection compares the MAC address, IP address, and source interface in all ARP packets to static entries in the ARP table. The following commands were introduced: **arp**, **arp-inspection**, and **show arp-inspection**.

Customizing the MAC Address Table for the Transparent Firewall

This section describes the MAC address table, and includes the following topics:

- Information About the MAC Address Table, page 6-11
- Licensing Requirements for ARP Inspection, page 6-12
- Default Settings, page 6-12
- Guidelines and Limitations, page 6-12
- Configuring the MAC Address Table, page 6-12
- Feature History for the MAC Address Table, page 6-14

Information About 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.

Licensing Requirements for ARP Inspection

The following table shows the licensing requirements for this feature.

Model	License Requirement
All models	Base License.

Default Settings

The default timeout value for dynamic MAC address table entries is 5 minutes.

By default, each interface automatically learns the MAC addresses of entering traffic, and the security appliance adds corresponding entries to the MAC address table.

Guidelines and Limitations

Context Mode Guidelines

- Supported in single and multiple context mode.
- In multiple context mode, configure the MAC address table within each context.

Firewall Mode Guidelines

Supported only in transparent firewall mode. Routed mode is not supported.

Configuring the MAC Address Table

This section describes how you can customize the MAC address table, and includes the following sections:

- Adding a Static MAC Address, page 6-13
- Disabling MAC Address Learning, page 6-13

Adding a Static MAC Address

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 "Adding a Static ARP Entry" section on page 6-9), a static MAC address entry is automatically added to the MAC address table.

To add a static MAC address to the MAC address table, perform the following steps:

Step 1 Choose the **Configuration** > **Device Setup** > **Bridging** > **MAC Address Table** pane.

Step 2 (Optional) To set the time a MAC address entry stays in the MAC address table before timing out, enter a value in the Dynamic Entry Timeout field.

This value is between 5 and 720 minutes (12 hours). 5 minutes is the default.

Step 3 Click **Add**.

The Add MAC Address Entry dialog box appears.

Step 4 From the Interface Name drop-down list, choose the source interface associated with the MAC address.

Step 5 In the MAC Address field, enter the MAC address.

Step 6 Click **OK**, and then **Apply**.

Disabling MAC Address Learning

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.

To disable MAC address learning, perform the following steps:

Step 1 Choose the **Configuration** > **Device Setup** > **Bridging** > **MAC Learning** pane.

Step 2 To disable MAC learning, choose an interface row, and click **Disable**.

Step 3 To reenable MAC learning, click **Enable**.

Step 4 Click **Apply**.

Feature History for the MAC Address Table

Table 6-2 lists the release history for this feature.

Table 6-4 Feature History for the MAC Address Table

Feature Name	Releases	Feature Information
MAC address table	7.0(1)	Transparent firewall mode uses a MAC address table. The following commands were introduced: **mac-address-table static**, **mac-address-table aging-time**, **mac-learn disable**, and **show mac-address-table**.

Firewall Mode Examples

This section includes examples of how traffic moves through the security appliance, and includes the following topics:

- How Data Moves Through the Security Appliance in Routed Firewall Mode, page 6-14
- How Data Moves Through the Transparent Firewall, page 6-20

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, page 6-15
- An Outside User Visits a Web Server on the DMZ, page 6-16
- An Inside User Visits a Web Server on the DMZ, page 6-17
- An Outside User Attempts to Access an Inside Host, page 6-18
- A DMZ User Attempts to Access an Inside Host, page 6-19

An Inside User Visits a Web Server

Figure 6-2 shows an inside user accessing an outside web server.

Figure 6-2 Inside to Outside

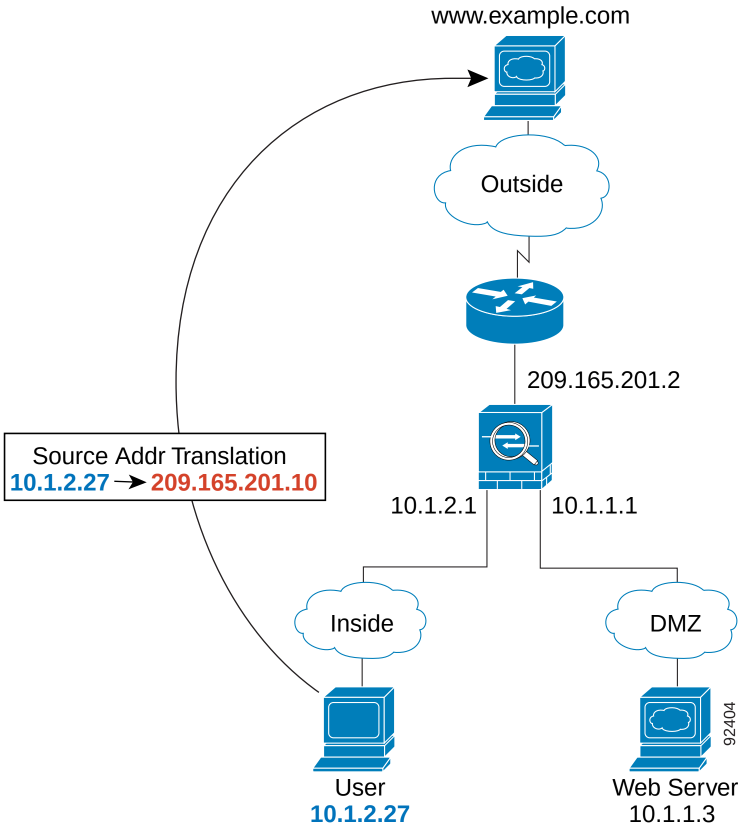

The following steps describe how data moves through the security appliance (see Figure 6-2):

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 6-3 shows an outside user accessing the DMZ web server.

Figure 6-3 Outside to DMZ

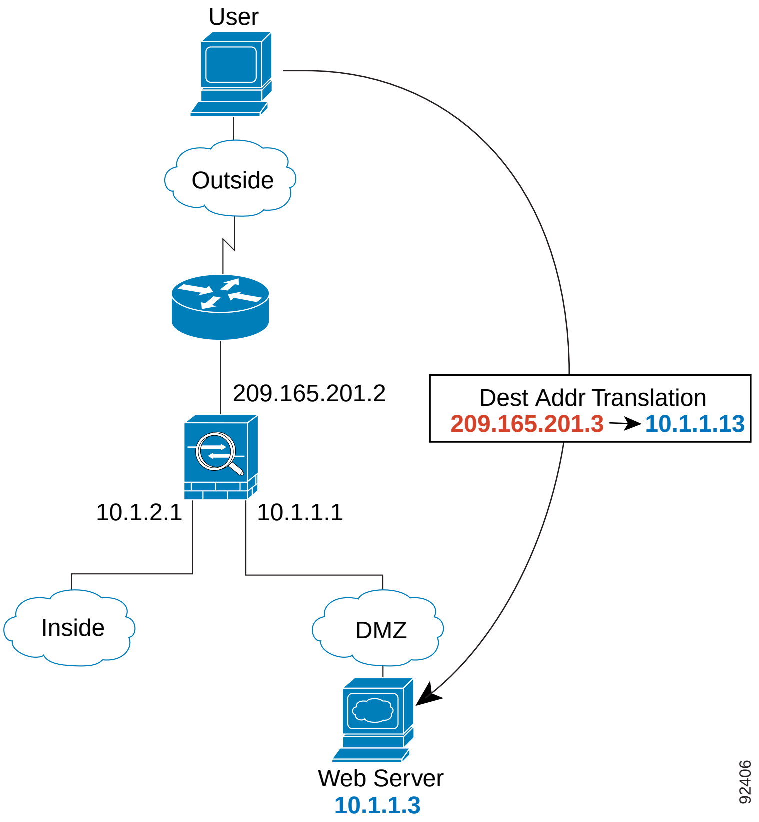

The following steps describe how data moves through the security appliance (see Figure 6-3):

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 6-4 shows an inside user accessing the DMZ web server.

Figure 6-4 Inside to DMZ

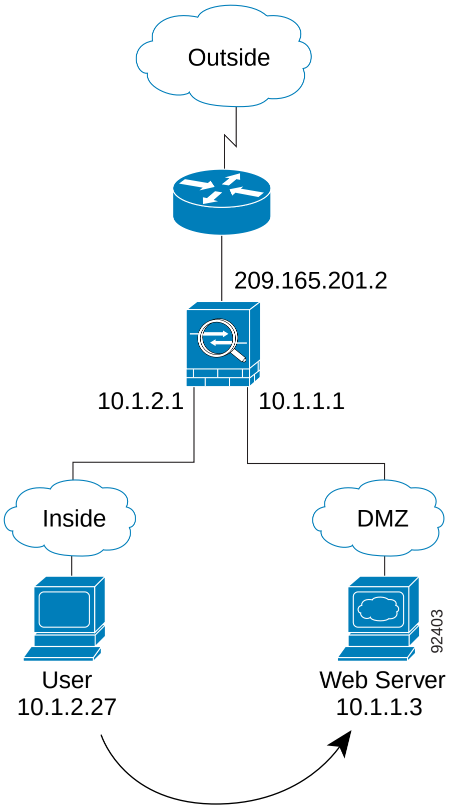

The following steps describe how data moves through the security appliance (see Figure 6-4):

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 6-5 shows an outside user attempting to access the inside network.

Figure 6-5 Outside to Inside

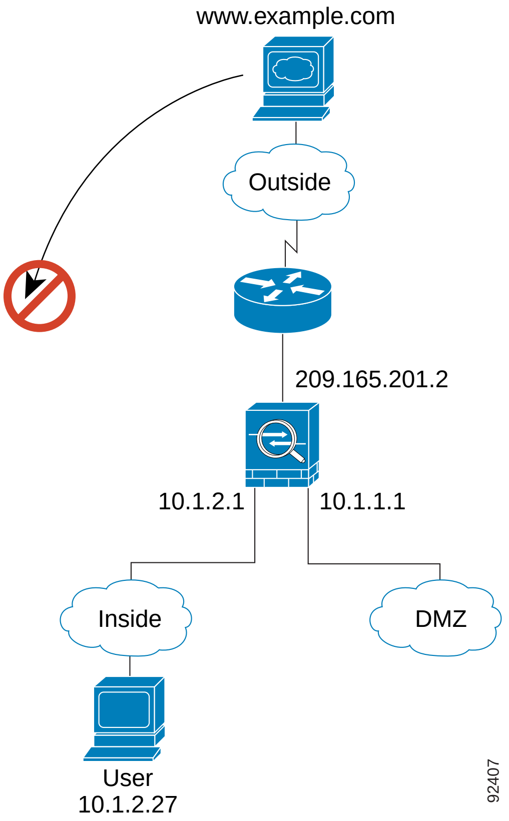

The following steps describe how data moves through the security appliance (see Figure 6-5):

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 6-6 shows a user in the DMZ attempting to access the inside network.

Figure 6-6 DMZ to Inside

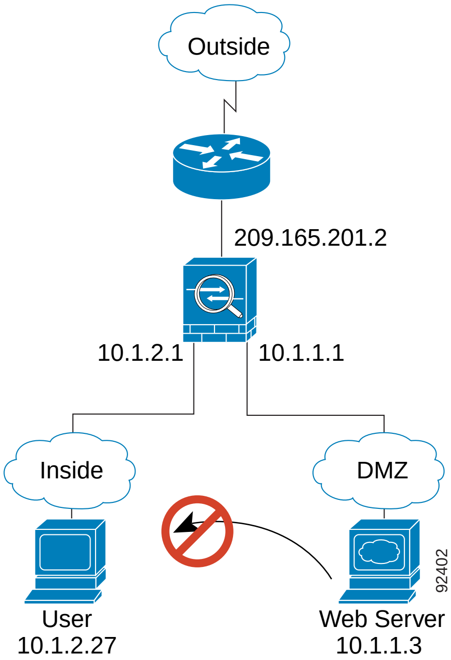

The following steps describe how data moves through the security appliance (see Figure 6-6):

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).

   The packet is denied, and the security appliance drops the packet and logs the connection attempt.

How Data Moves Through the Transparent Firewall

Figure 6-7 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 6-7 Typical Transparent Firewall Data Path

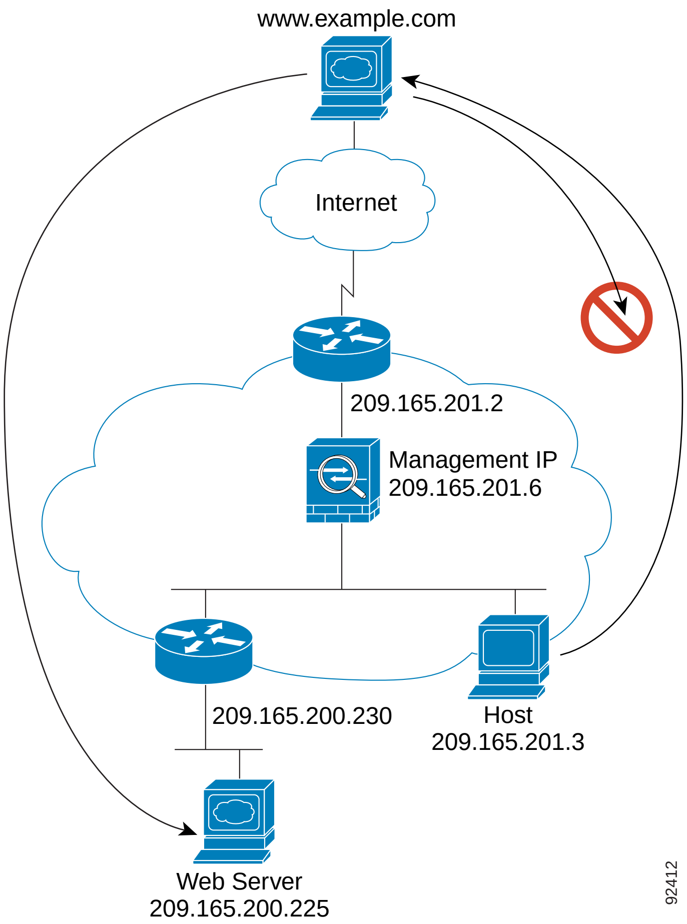

This section describes how data moves through the security appliance, and includes the following topics:

- An Inside User Visits a Web Server, page 6-21
- An Inside User Visits a Web Server Using NAT, page 6-22
- An Outside User Visits a Web Server on the Inside Network, page 6-23
- An Outside User Attempts to Access an Inside Host, page 6-24

An Inside User Visits a Web Server

Figure 6-8 shows an inside user accessing an outside web server.

Figure 6-8 Inside to Outside

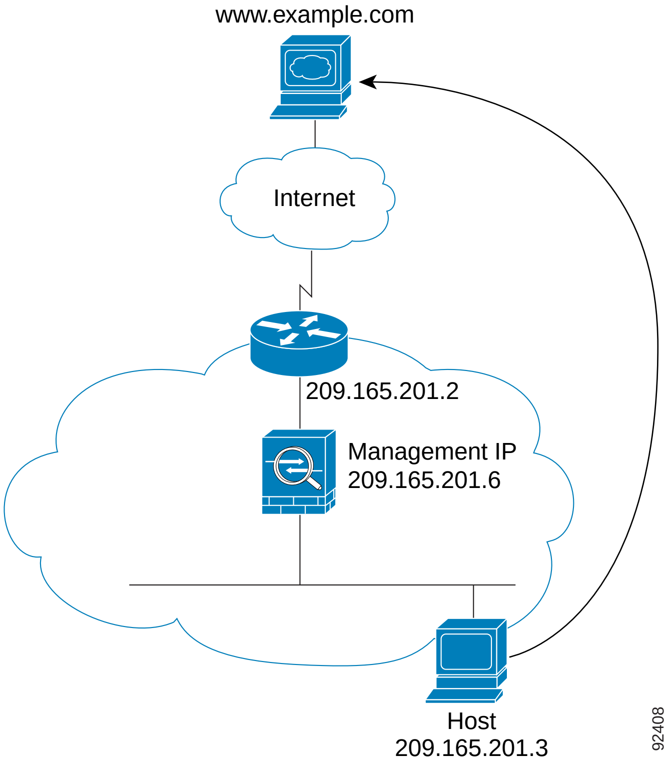

The following steps describe how data moves through the security appliance (see Figure 6-8):

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 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. The web server responds to the request; 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 Inside User Visits a Web Server Using NAT

Figure 6-8 shows an inside user accessing an outside web server.

Figure 6-9 Inside to Outside with NAT

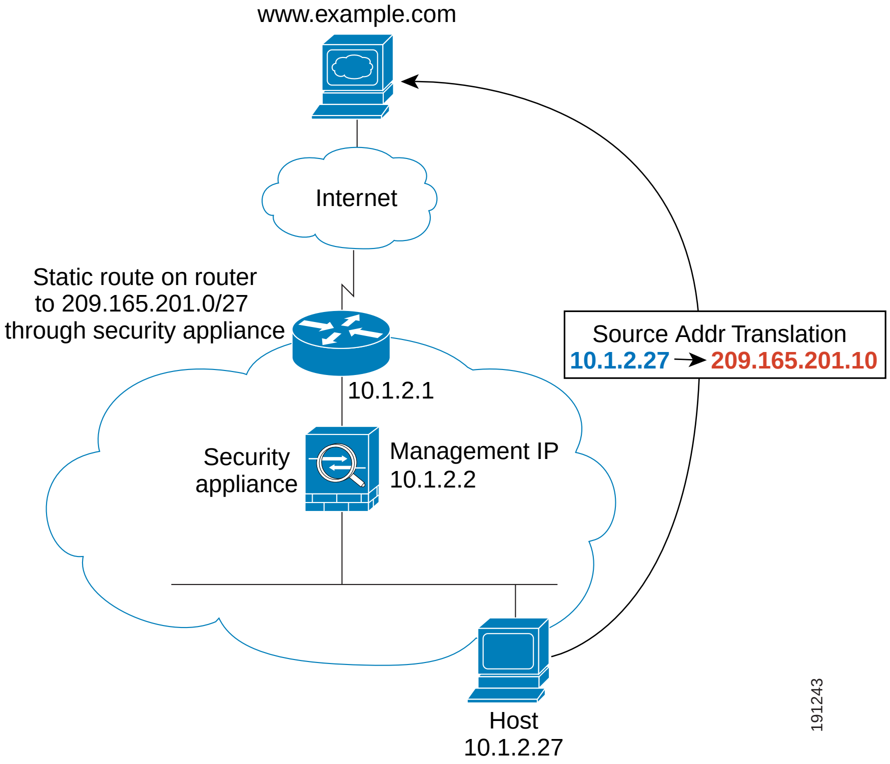

The following steps describe how data moves through the security appliance (see Figure 6-8):

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 translates the real address (10.1.2.27) to the mapped address 209.165.201.10.

   Because the mapped address is not on the same network as the outside interface, then be sure the upstream router has a static route to the mapped network that points to the security appliance.
4. The security appliance then records that a session is established and forwards the packet from the outside interface.
5. 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, 209.165.201.2.

   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.
6. The web server responds to the request; because the session is already established, the packet bypasses the many lookups associated with a new connection.

7. The security appliance performs NAT by translating the mapped address to the real address, 10.1.2.27.

An Outside User Visits a Web Server on the Inside Network

Figure 6-10 shows an outside user accessing the inside web server.

Figure 6-10 Outside to Inside

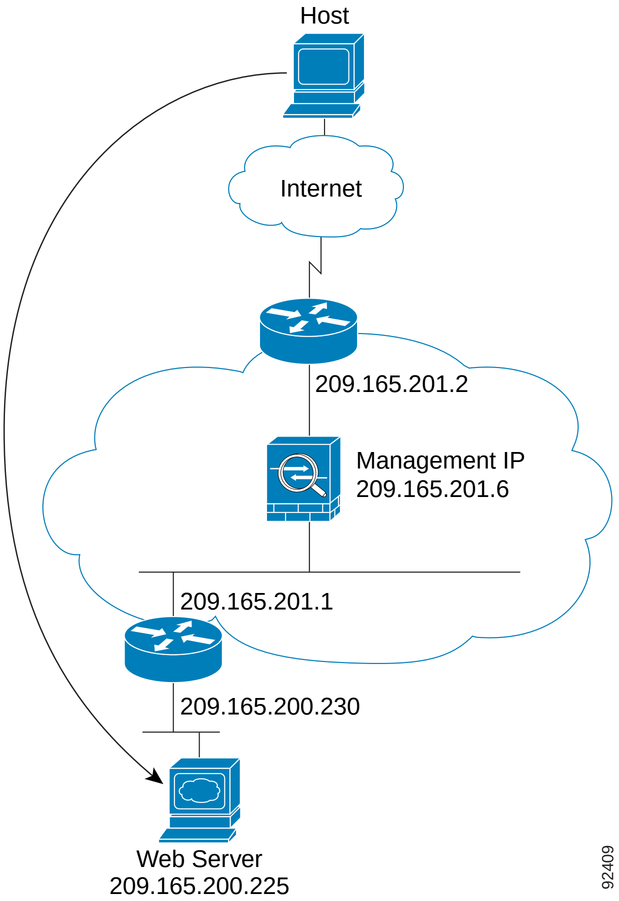

The following steps describe how data moves through the security appliance (see Figure 6-10):

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. The web server responds to the request; 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 6-11 shows an outside user attempting to access a host on the inside network.

Figure 6-11 Outside to Inside

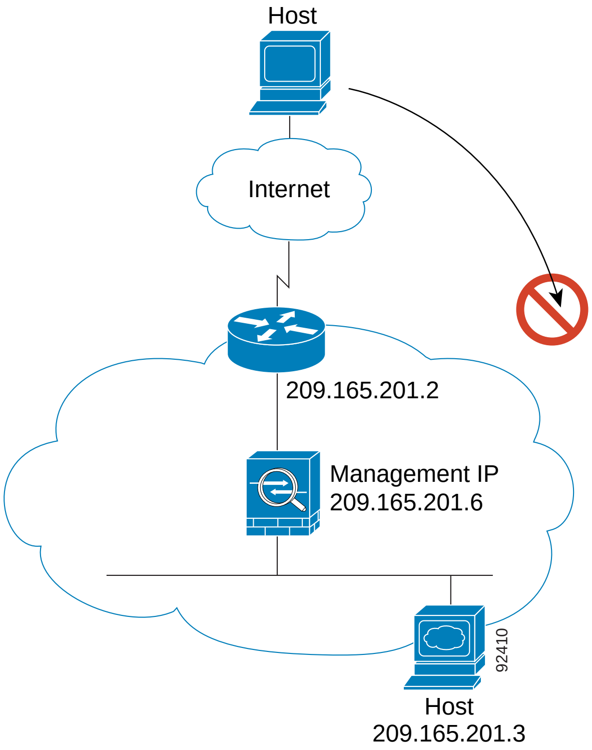

The following steps describe how data moves through the security appliance (see Figure 6-11):

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.

PART 2

Device Setup and Management

CHAPTER 7

Using the Startup Wizard

The ASDM Startup Wizard guides you through the initial configuration of the security appliance, and helps you define its settings. This chapter includes the following sections:

Information About the Startup Wizard

To access this feature in the main ASDM application window, choose one of the following:

- **Wizards** > **Startup Wizard**.
- **Configuration** > **Device Setup** > **Startup Wizard,** and then click **Launch Startup Wizard**.

For more information, see the "Starting ASDM from a Web Browser" section on page 1-7, the *Cisco ASA 5500 Series Adaptive Security Appliance Getting Started Guide*, and the *Cisco ASA 5505 Getting Started Guide*.

Licensing Requirements for the Startup Wizard

The following table shows the licensing requirements for this feature:

Model	License Requirement
All models	Base License

Prerequisites for the Startup Wizard

To complete the Startup Wizard, make sure that you have the following information available:

- The hostname
- The domain name
- A password to restrict administrative access through ASDM or the CLI
- The IP address information of the outside interface
- Other interfaces, such as the inside or DMZ interfaces
- NAT or PAT rules
- DHCP settings for the inside interface, for use with a DHCP server

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context modes, as noted in Table 7-1.

Firewall Mode Guidelines

Supported in routed and transparent firewall modes, as noted in Table 7-1.

Failover Guidelines

Supports sessions in Stateful Failover.

IPv6 Guidelines

Supports IPv6.

Model Guidelines

Supports all models.

Startup Wizard Screens for ASA 5500 Series Adaptive Security Appliances

Table 7-1 lists all of the required Startup Wizard screens for configuring the ASA 5500 series security appliances. The actual sequence of screens is determined by your specified configuration selections. The Availability columns lists the mode or modes in which each screen appears and provides additional configuration information. Click the name to view information for the selected screen.

Table 7-1 Startup Wizard Screens for ASA 5500 Series Adaptive Security Appliances

Screen Name	Availability
Step 1 - Starting Point or Welcome, page 7-4	All modes.
Step 2 - Basic Configuration, page 7-5	

Table 7-1 Startup Wizard Screens for ASA 5500 Series Adaptive Security Appliances

Screen Name	Availability
Step 3 - Auto Update Server, page 7-5	Single, routed and transparent modes. If enabled in single transparent mode, the Interface Configuration and Step 13 - DHCP Server screens are not available.
Step 4 - Management IP Address Configuration, page 7-6	Single, transparent mode only.
Outside Interface Configuration, page 7-28	Single, routed mode only.
Outside Interface Configuration - PPPoE, page 7-27	
Interface Configuration, page 7-27	Single, transparent mode only.
Other Interfaces Configuration, page 7-16	All modes.
Step 12 - Static Routes, page 7-11	
Step 13 - DHCP Server, page 7-11	
Step 14 - Address Translation (NAT/PAT), page 7-12	Single, routed mode only.
Step 15 - Administrative Access, page 7-13	All modes.
Step 17 - Startup Wizard Summary, page 7-15	

Startup Wizard Screens for the ASA 5505 Adaptive Security Appliance

Table 7-2 lists all of the required Startup Wizard screens for configuring the ASA 5505 security appliance only. The sequence of screens listed represents configuration for the single, routed mode. The Availability columns lists the mode or modes in which each screen appears and provides additional configuration information. Click the name to view information for the selected screen.

Table 7-2 Startup Wizard Screens for the ASA 5505 Adaptive Security Appliance

Screen Name and Sequence	Availability
Step 1 - Starting Point or Welcome, page 7-4	All modes. The Teleworker option in Step 2 - Basic Configuration is available only on the ASA 5505.
Step 2 - Basic Configuration, page 7-5	
Step 3 - Auto Update Server, page 7-5	Single, routed and transparent modes. Enabled only if configured for teleworker usage.
Step 4 - Management IP Address Configuration, page 7-6	Single, transparent mode only.

Table 7-2 Startup Wizard Screens for the ASA 5505 Adaptive Security Appliance (continued)

Screen Name and Sequence	Availability
Step 5 - Interface Selection, page 7-6	Single, routed mode only.
Step 6 - Switch Port Allocation, page 7-7	
Step 7 - Interface IP Address Configuration, page 7-7	
Step 8 - Internet Interface Configuration - PPPoE, page 7-8	
Step 9 - Business Interface Configuration - PPPoE, page 7-9	
Step 10 - Home Interface Configuration - PPPoE, page 7-10	
Step 11 - General Interface Configuration, page 7-10	
Step 12 - Static Routes, page 7-11	All modes. Enabled only if configured for teleworker usage.
Step 13 - DHCP Server, page 7-11	All modes.
Step 14 - Address Translation (NAT/PAT), page 7-12	Single, routed mode only.
Step 15 - Administrative Access, page 7-13	All modes.
Step 16 - Easy VPN Remote Configuration, page 7-14	Single, routed mode, only when enabled for teleworker usage.
Step 17 - Startup Wizard Summary, page 7-15	All modes.

Step 1 - Starting Point or Welcome

Step 1 To change the existing configuration, click the **Modify existing configuration** radio button.

Step 2 To set the configuration at the factory default values for the inside interface, click the **Reset configuration to factory defaults** radio button.

Step 3 To configure the IP address and subnet mask of the management interface, check the **Configure the IP address of the management interface** check box.

Step 4 Specify the IP address of the management interface.

Step 5 Choose the subnet mask of the management interface from the drop-down list.

Note If you reset the configuration to factory defaults, you cannot undo these changes by clicking **Cancel** or by closing this screen.

Step 6 Click **Next** to continue.

Step 2 - Basic Configuration

Step 1 To specify a group of configuration settings for a remote worker, check the **Configure the device for Teleworker usage** check box. For more information, see Step 16 - Easy VPN Remote Configuration, page 7-14.

Step 2 Specify a hostname for the security appliance. The hostname can be up to 63 alphanumeric characters in mixed case.

Step 3 Specify the IPSec domain name of the security appliance, which can be used for certificates. The domain name can be a maximum of 63 alphanumeric characters, with no special characters or spaces.

Step 4 The privileged mode (enable) password is required to administer the security appliance through ASDM or the CLI. To change the current privileged mode (enable) password, check the **Change privileged mode (enable) password** check box.

Note If you leave the password field blank, a Password Confirmation dialog box appears to notify you that to do so is a high security risk.

Step 5 Specify the old enable password, if one exists.

Step 6 Specify the new enable password. The password is case-sensitive and can be up to 32 alphanumeric characters.

Step 7 Reenter the new enable password.

Step 8 Click **Next** to continue.

Step 3 - Auto Update Server

Step 1 To enable communication between the security appliance and an Auto Update Server, check the **Enable Auto Update** check box.

Step 2 To define the beginning of the URL for the Auto Update Server, from the Server URL drop-down list, choose either HTTPS or HTTP.

Step 3 To confirm that an SSL certificate is enabled on the Auto Update Server, check the **Verify Server SSL certificate** check box.

Step 4 Enter the username to log in to the Auto Update Server.

Step 5 Enter the password to log in to the Auto Update Server.

Step 6 Reenter the password to confirm it.

Step 7 To uniquely identify the security appliance, choose the Device ID Type from the drop-down list. To enable the Device ID field and specify a particular name, choose User-defined name.

Step 8 Enter a unique string to use as the security appliance ID.

Step 9 Click **Next** to continue.

Step 4 - Management IP Address Configuration

Step 1 Specify the IP address of the host that can access this context for management purposes using ASDM or a session protocol.

Step 2 Specify the subnet mask for the management IP address.

Step 3 Click **Next** to continue.

Step 5 - Interface Selection

This screen allows you to group the eight, Fast Ethernet switch ports on the ASA 5505 into three VLANs. These VLANs function as separate, Layer 3 networks. You can then choose or create the VLANs that define your network—one for each interface: outside (Internet), inside (Business), or DMZ (Home). A DMZ is a separate network located in the neutral zone between a private (inside) network and a public (outside) network.

To create three VLANs to define your network, perform the following steps:

Step 1 In the Outside VLAN or Internet VLAN area, do the following:

a. From the Choose a VLAN drop-down list, choose a predefined outside VLAN by number.

b. To create a new outside VLAN, check the **Create a VLAN** check box.

c. To enable the outside VLAN, check the **Enable VLAN** check box.

Step 2 In the Inside VLAN or Business VLAN area, do the following:

a. From the Choose a VLAN drop-down list, choose a predefined inside VLAN by number.

b. To create a new inside VLAN, check the **Create a VLAN** check box.

c. To enable the inside VLAN, check the **Enable VLAN** check box.

Step 3 In the DMZ VLAN or Home VLAN (Optional) area, do the following:

a. From the Choose a VLAN drop-down list, choose a predefined inside VLAN by number.

b. To create a new inside VLAN, check the **Create a VLAN** check box.

- To disable configuration of this VLAN, check the **Do not configure** check box.

Step 4 Click **Next** to continue.

Step 6 - Switch Port Allocation

This screen lets you allocate switch ports to outside (Internet), inside (Business), or DMZ (Home) interfaces. The DMZ interface is not available in transparent mode. You must add the ports to the associated VLANs. By default, all switch ports begin with VLAN1.

Step 1 In the Switch Ports for Outside VLAN (vlan*id*) or Switch Ports for Internet VLAN (vlan*id*) area, do the following:

a. Choose a port to add or remove from the available list of ports.

b. Choose a port to add or remove from the allocated list of ports.

c. To add a port to the available or allocated list of ports, click **Add**.

d. To remove a port from the available or allocated list of ports, click **Remove**.

Step 2 In the Switch Ports for Inside VLAN (vlan*id*) or Switch Ports for Business VLAN (vlan*id)* area, do the following:

a. Choose a port to add or remove from the available list of ports.

b. Choose a port to add or remove from the allocated list of ports.

c. To add a port to the available or allocated list of ports, click **Add**.

d. To remove a port from the available or allocated list of ports, click **Remove**.

Step 3 In the Switch Ports for DMZ VLAN (vlan*id*) or Switch Ports for Home VLAN (vlan*id)* area, do the following:

a. Choose a port to add or remove from the available list of ports.

b. Choose a port to add or remove from the allocated list of ports.

c. To add a port to the available or allocated list of ports, click **Add**.

d. To remove a port from the available or allocated list of ports, click **Remove**.

Step 4 Click **Next** to continue.

Step 7 - Interface IP Address Configuration

To configure the interface through a PPPoE server, a DHCP server, or by specifying a particular IP address and subnet mask, perform the following steps:

Step 1 In the Outside IP Address or Internet IP Address area, do one of the following:

- To specify an outside IP address, click the **Use the following IP address** radio button.
- Enter the specific outside IP address and choose the subnet mask from the drop-down list.
- To obtain an outside IP address from a DHCP server, click the **Use DHCP** radio button.
- To obtain the default route for an outside IP address from a DHCP server, check the **Obtain default rote using DHCP** check box.
- To obtain an outside IP address from a PPPoE server, click the **Use PPPoE** radio button.

Step 2 In the Inside IP Address or Business IP Address area, do one of the following:

- To specify an inside IP address, click the **Use the following IP address** radio button.
- Enter the specific inside IP address and choose the subnet mask from the drop-down list.
- To obtain an inside IP address from a DHCP server, click the **Use DHCP** radio button.
- To obtain the default route for an inside IP address from a DHCP server, check the **Obtain default rote using DHCP** check box.
- To obtain an inside IP address from a PPPoE server, click the **Use PPPoE** radio button.

In the DMZ IP Address or Home IP Address area, choose one of the following:

- To specify a DMZ IP address, click the **Use the following IP address** radio button, then enter the specific DMZ IP address and choose the subnet mask from the drop-down list.
- To obtain a DMZ IP address from a DHCP server, click the **Use DHCP** radio button.
- To obtain a DMZ IP address from a PPPoE server, click the **Use PPPoE** radio button.

Step 3 Click **Next** to continue.

Step 8 - Internet Interface Configuration - PPPoE

Note For all ASA 5500 series models except ASA 5505, with a full license, the security appliance supports up to five interfaces, with a maximum of three outside interfaces. In restricted mode, the security appliance supports up to three interfaces, and in transparent mode, the security appliance supports up to two interfaces. After you have created the maximum number of interfaces, or the maximum number of interfaces has already been named, you may not be able to create a new VLAN, and must select an existing one.

Step 1 Specify the name of the group on the PPPoE server. You must specify a group name to proceed.

Step 2 In the User Authentication area, do the following:

a. Specify your username on the PPPoE server.

b. Specify your password on the PPPoE server.

c. Confirm the PPPoE password that you entered.

Step 3 In the Authentication Method area, do one of the following:

- To use PAP authentication, click the **PAP** radio button.
- To use CHAP authentication, click the **CHAP** radio button.
- To use MS-CHAP authentication, click the **MS-CHAP** radio button.

Step 4 In the IP Address area, do one of the following:

- To obtain an IP address for the interface from the PPPoE server, click the **Obtain an IP address using PPPoE** radio button. This field is not visible in transparent mode.
- Specify an IP address for the Internet interface. This field is not visible in transparent mode.
  - Specifies the IP address that you want to use for the Internet interface.
  - Choose a subnet mask for the Internet interface from the drop-down list.

- To set the default routing using the PPPoE server, check the **Obtain default route using PPPoE** check box.

Step 5 Click **Next** to continue.

Step 9 - Business Interface Configuration - PPPoE

Note For all ASA 5500 series models except ASA 5505, with a full license, the security appliance supports up to five interfaces, with a maximum of three outside interfaces. In restricted mode, the security appliance supports up to three interfaces, and in transparent mode, the security appliance supports up to two interfaces. After you have created the maximum number of interfaces, or the maximum number of interfaces has already been named, you may not be able to create a new VLAN, and must select an existing one.

Step 1 Enter the name of the group on the PPPoE server. You must specify a group name to proceed.

Step 2 In the User Authentication area, do the following:

a. Enter your username on the PPPoE server.

b. Enter your password on the PPPoE server.

c. Enter the PPPoE password that you entered.

Step 3 In the Authentication Method area, choose one of the following:

- To use PAP authentication, click **PAP**.
- To use CHAP authentication., click **CHAP**.
- To use MS-CHAP authentication, click **MS-CHAP**.

Step 4 In the IP Address area, choose one of the following:

- Click the **Obtain an IP address using PPPoE** radio button to obtain an IP address for the interface from the PPPoE server. This option is not visible in transparent mode.
- Enter an IP address for the inside interface. This option is not visible in transparent mode.
  - Enter the IP address that you want to use for the inside interface.
  - Choose a subnet mask for the Internet interface from the drop-down list.
- To set the default routing using the PPPoE server, check the **Obtain default route using PPPoE** check box.

Step 5 Click **Next** to continue.

Step 10 - Home Interface Configuration - PPPoE

Note For all ASA 5500 series models except ASA 5505, with a full license, the security appliance supports up to five interfaces, with a maximum of three outside interfaces. In restricted mode, the security appliance supports up to three interfaces, and in transparent mode, the security appliance supports up to two interfaces. After you have created the maximum number of interfaces, or the maximum number of interfaces has already been named, you may not be able to create a new VLAN, and must select an existing one.

Step 1 Enter the name of the group on the PPPoE server. You must specify a group name to proceed.

Step 2 In the User Authentication area, do the following:

a. Enter your username on the PPPoE server.

b. Enter your password on the PPPoE server.

c. Enter the PPPoE password that you entered.

Step 3 In the Authentication Method area, choose one of the following:

- To use PAP authentication, click **PAP**.
- To use CHAP authentication., click **CHAP**.
- To use MS-CHAP authentication, click **MS-CHAP**.

Step 4 In the IP Address area, choose one of the following:

- Click the **Obtain an IP address using PPPoE** radio button to obtain an IP address for the DMZ interface from the PPPoE server. This option is not visible in transparent mode.
- Enter an IP address for the DMZ interface. This option is not visible in transparent mode.
  - Enter the IP address that you want to use for the DMZ interface.
  - Choose a subnet mask for the DMZ interface from the drop-down list.
- To set the default routing using the PPPoE server, check the **Obtain default route using PPPoE** check box.

Step 5 Click **Next** to continue.

Step 11 - General Interface Configuration

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

To enable and restrict traffic between interfaces and between hosts connected to the same interface, perform the following steps:

Step 1 To enable traffic between two or more interfaces with the same security level, check the **Enable traffic between two or more interfaces with the same security level** check box.

Step 2 To enable traffic between two or more hosts connected to the same interface, check the **Enable traffic between two or more hosts connected to the same interface** check box.

Step 3 In the Restrict traffic area, do the following:

- To restrict traffic from an interface, choose an interface from the drop-down list.
- To restrict traffic to an interface, choose an interface from the drop-down list.

Step 4 Click **Next** to continue.

Step 12 - Static Routes

Step 1 To create, edit, and remove static routes that will access networks connected to a router on any interface, perform the following steps:

Step 2 Choose to filter by IPv4 addresses, IPv6 addresses, or both.

Step 3 To continue, see the following sections:

- Static Routes, page 12-40
- Adding or Editing Static Routes, page 7-11

Step 4 Click **Next** to continue.

Adding or Editing Static Routes

The Add/Edit Static Routes dialog box lets you add, edit, or remove a static route. For more information, see the "Static Routes" section on page 12-40.

Step 13 - DHCP Server

Step 1 To allow connection to the DHCP server from the inside interface, check the **Enable DHCP server on the inside interface** check box.

Step 2 In the DHCP Address Pool area, do the following:

- Enter the starting range of the DHCP server pool in a block of IP addresses from the lowest to highest.
- Enter the ending range of the DHCP server pool in a block of IP addresses from the lowest to highest.

Note The security appliance supports up to 256 IP addresses.

Step 3 In the DHCP Parameters area, do the following:

a. To allow automatic configuration of the DNS server, WINS server, lease length, and ping timeout settings, check the **Enable auto-configuration** check box.

b. Enter the IP address of the DNS server.

c. Enter the IP address of the WINS server.

d. Enter the IP address of the alternate DNS server.

e. Enter the IP address of the alternate WINS server.

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

g. Enter the parameters for the ping timeout value in milliseconds.

h. Enter the domain name of the DNS server to use DNS.

i. To enable DHCP auto-configuration and choose the interface from the drop-down list, check the **Enable auto-configuration from interface** check box. The values you specify in the previous areas of this screen take precedence over the auto-configured values.

Step 4 Click **Next** to continue.

For more information, see the "DHCP Server" section on page 15-4.

Step 14 - Address Translation (NAT/PAT)

PAT lets you set up a single IP address for use as the global address. In addition, you can set multiple outbound sessions to appear as if they originate from a single IP address. 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 all 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.

When you use PAT, be aware of the following:

- PAT 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.
- PAT does not work with the **established** command.
- With passive FTP, use the **inspect protocol ftp strict** command with the **access-list** command to allow outbound FTP traffic.
- A DNS server on a higher level security interface cannot use PAT.

Step 1 To enable NAT and share several external IP addresses on the inside VLAN to be used for translation, click the **Use Network Address Translation (NAT)** radio button, then do the following:

a. Enter the first IP address in a range of IP addresses to be used for translation.

b. Enter the last IP address in a range of IP addresses to be used for translation.

c. (Optional) Enter the subnet mask for the range of IP addresses to be used for translation.

Step 2 To enable PAT, click the **Use Port Address Translation (PAT)** radio button. If you select this option, choose one of the following:

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

- To use the IP address of the outside interface for PAT, click the **Use the IP address on the outside interface** radio button.

- To indicate a particular address to use for PAT, click the **Specify an IP address** radio button.
  - Enter an IP address for the outside interface for PAT.
  - (Optional) Choose a subnet mask from the drop-down list.
- To allow traffic through the security appliance without translation, check the **Enable traffic through the firewall without translation** check box.

Step 3 Click **Next** to continue.

Step 15 - Administrative Access

To configure management access on the security appliance, perform the following steps:

Step 1 To add or change the access type, an interface, and then specify the IP address and netmask of the host network that may connect to that interface for management purposes only, see the "Adding or Editing Administrative Access Entry" section on page 7-13.

- The Type column specifies whether the host or network is accessing the security appliance through HTTP over SSL in ASDM, SSH, or Telnet.
- The Interface column displays the host or network name.
- The IP Address column displays the IP address of the host or network.
- The Mask column displays the subnet mask of the host or network.

Step 2 To enable a secure connection to an HTTP server to access ASDM, check the **Enable HTTP server for HTTPS/ASDM access** check box.

Step 3 To allow ASDM to collect and display statistics, check the **Enable ASDM history metrics** check box.

Step 4 Click **Next** to continue.

Adding or Editing Administrative Access Entry

Step 1 To configure the hosts, in the main ASDM application window, choose one of the following:

- **Configuration** > **Properties** > **Device Access** > **HTTPS/ASDM**
- **Configuration** > **Properties** > **Device Access** > **Telnet**
- **Configuration** > **Properties** > **Device Access** > **SSH**
- **Configuration** > **Properties** > **History Metrics**

Step 2 Choose one of the following preconfigured connections for the CLI console sessions from the Access Type drop-down list:

- ASDM/HTTPS
- SSH
- Telnet

Note ASDM uses HTTP over SSL (HTTPS) for all communication with the security appliance.

Step 3 Choose the interface name from the Interface drop-down list.

Step 4 Enter an IP address for the interface.

Step 5 Enter a subnet mask for the interface from the Subnet Mask drop-down list.

Step 6 Click **OK** to save these settings and return to the Administrative Access screen.

Step 16 - Easy VPN Remote Configuration

The security appliance can act as an Easy VPN remote device to enable deployment of VPNs to remote locations. The following two modes of operation are available:

- Client Mode
- Network Extension Mode

In Client Mode, the security appliance does not expose the IP addresses of clients on the inside network. Instead, the security appliance uses NAT to translate the IP addresses on the private network to a single, assigned IP address. In this mode, you cannot ping or access any device from outside the private network.

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

To configure the security appliance in one of these two modes, use the following guidelines:

- 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.

Note To access this screen, you must have checked the **Configure the device for Teleworker usage** check box in Step 2 - Basic Configuration and unchecked the **Enable Auto Update** check box in the Interface Configuration.

To form a secure VPN tunnel between the security appliance and a remote Cisco VPN 3000 concentrator, Cisco router, or security appliance that is acting as an Easy VPN server, perform the following steps:

Step 1 To enable the security appliance to act as an Easy VPN remote device, check the **Enable Easy VPN remote** check box. If you do not enable this feature, any host that has access to the security appliance outside interface through a VPN tunnel can manage it remotely.

Step 2 In the Mode area, choose one of the following:

- If you are using a DHCP server to generate dynamic IP addresses for hosts on your inside network, click the **Client mode** radio button.
- If hosts on your inside network have static IP addresses, click the **Network Extension mode** radio button.

Step 3 In the Group Settings area, do the following:

a. To use X.509 certificates to enable the IPSec main mode, click the **Use X.509 Certificate** radio button. Choose the trustpoint from the drop-down list.

b. To enter a password for a group of users, click the **Use group password** radio button.

- Enter a name for the user group.
- Enter a password for the user group.
- Confirm the password.

Step 4 In the User Settings area, do the following:

a. Enter a username for your settings.

b. Enter a password for your settings.

c. Confirm the password for your settings.

Step 5 In the Easy VPN Server area, do the following:

a. Enter the IP address of the primary Easy VPN server.

b. Enter the IP address of a secondary Easy VPN server.

Note The security appliance supports a maximum of 11 Easy VPN servers: one primary and up to ten secondary. 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 ISP. After you have connected the ASA 5500 series security appliance to the DSL or cable modem, follow the instructions provided by your ISP to complete the network connection. You can obtain an IP address through a PPPoE server, a DHCP server, or a static configuration.

Step 6 Click **Next** to continue.

Step 17 - Startup Wizard Summary

This screen summarizes all of the configuration settings that you have made for the security appliance.

Step 1 To change any of the settings in previous screens, click **Back**.

Step 2 Choose one of the following:

- If you ran the Startup Wizard directly from a browser, when you click **Finish**, the configuration settings that you created through the wizard are sent to the security appliance and saved in flash memory automatically.
- If you ran the Startup Wizard from within ASDM, you must explicitly save the configuration in flash memory by choosing **File > Save Running Configuration to Flash**.

Other Interfaces Configuration

To configure the remaining interfaces, perform the following steps:

Step 1 Select the interface to change and click **Edit**.

The Edit Interface dialog box appears.

- The Interface field displays the network interface on which the original host or network resides.
- The Name field displays the name of the interface being configured.
- The Security Level field 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.

Step 2 To assign the same security level to two or more interfaces, and enable traffic between them, check the **Enable traffic between two or more interfaces with same security levels** check box.

Step 3 If you have an interface between two or more hosts and want to enable traffic between them, check the **Enable traffic between two or more hosts connected to the same interface** check box.

Step 4 Click **Next** to continue.

Editing Interfaces

On the Interface Properties and IPv4 Settings tab, perform the following steps:

Step 1 The Interface field is display-only and shows the name of the selected interface to edit. The Enable interface check box is checked by default.

Step 2 The Interface Name field displays the name of the selected interface. Change the name of the interface, if needed.

Step 3 The Security Level field displays the security level of the selected interface. Change the security level for the interface, if needed. If you change the security level of the interface to a lower level, a warning message appears.

Step 4 In the IP Address area, choose one of the following three options:

- To enter a specific IP address for an interface, click the **Use the following IP address** radio button.
  - Enter the IP address of the interface.
  - Choose an existing subnet mask from the drop-down list.
- To use the security appliance as a DHCP server, click the **Use DHCP** radio button.
- To use PPPoE to provide an authenticated method of assigning an IP address to an outside interface, click the **Use PPPoE** radio button.

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

  - Enter a group name to proceed.
  - Enter the PPPoE username and password, and confirm the password.

- PAP is the default authentication method for PPPoE. You have the option of configuring CHAP or MS-CHAP manually by clicking the applicable radio button.

Step 5 Choose one of the following:

- To assign an IP address using PPPoE, click the **Obtain IP Address using PPPoE** radio button.
- To assign a particular IP address, click the **Specify an IP address** radio button.
  - Enter the IP address.
  - Choose a subnet mask from the drop-down list.

Step 6 Click **OK** to save these settings.

On the IPv6 Settings tab, perform the following steps:

Step 1 To configure Neighbor Discovery settings, see the "Configuring IPv6 Neighbor Discovery" section on page 7-17.

Step 2 To configure IPv6 addresses on an interface, see the "Configuring IPv6 Addresses on an Interface" section on page 7-20.

Step 3 To configure IPv6 prefixes on an interface, see the "Configuring IPv6 Prefixes on an Interface" section on page 7-21.

Step 4 Click **OK** to save these settings.

Configuring IPv6 Neighbor Discovery

The IPv6 neighbor discovery process uses ICMPv6 messages and solicited-node multicast addresses to determine the link-layer address of a neighbor on the same network (local link), verify the reachability of a neighbor, and keep track of neighboring routers. For more information about IPv6 Neighbor Discovery, see the "Information About IPv6 Neighbor Discovery" section on page 14-1.

This section includes the following topics:

- Configuring Neighbor Solicitation Messages, page 7-17
- Configuring Router Advertisement Messages, page 7-22
- Configuring IPv6 Static Neighbors, page 7-25

Configuring Neighbor Solicitation Messages

Neighbor solicitation messages (ICMPv6 Type 135) are sent on the local link by nodes attempting to discover the link-layer addresses of other nodes on the local link. The neighbor solicitation message is sent to the solicited-node multicast address. The source address in the neighbor solicitation message is the IPv6 address of the node sending the neighbor solicitation message. The neighbor solicitation message also includes the link-layer address of the source node.

After receiving a neighbor solicitation message, the destination node replies by sending a neighbor advertisement message (ICPMv6 Type 136) on the local link. The source address in the neighbor advertisement message is the IPv6 address of the node sending the neighbor advertisement message; the destination address is the IPv6 address of the node that sent the neighbor solicitation message. The data portion of the neighbor advertisement message includes the link-layer address of the node sending the neighbor advertisement message. After the source node receives the neighbor advertisement, the source node and destination node can communicate.

Figure 7-1 shows the neighbor solicitation and response process.

Figure 7-1 IPv6 Neighbor Discovery—Neighbor Solicitation Message

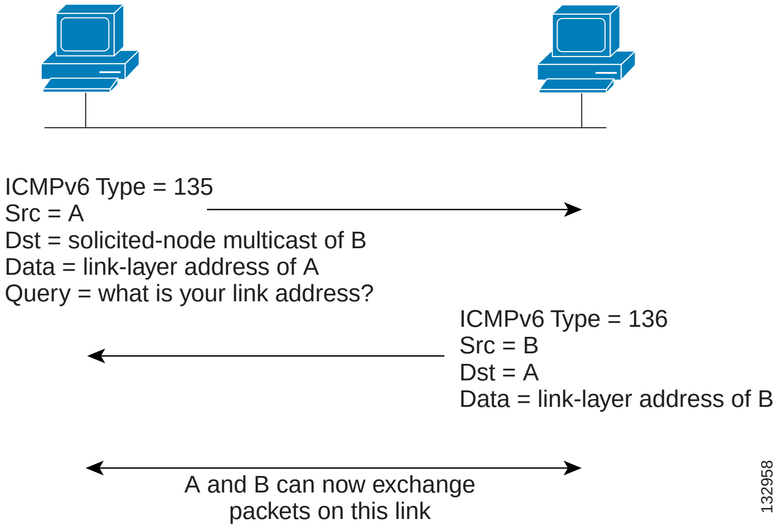
ICMPv6 Type = 135
Src = A
Dst = solicited-node multicast of B
Data = link-layer address of A
Query = what is your link address?
ICMPv6 Type = 136
Src = B
Dst = A
Data = link-layer address of B
A and B can now exchange
packets on this link
132958

Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. When a node wants to verify the reachability of a neighbor, the destination address in a neighbor solicitation message is the unicast address of the neighbor.

Neighbor advertisement messages are also sent when there is a change in the link-layer address of a node on a local link. When there is such a change, the destination address for the neighbor advertisement is the all-nodes multicast address.

You can configure the neighbor solicitation message interval and neighbor reachable time on a per-interface basis.

In addition, you can configure DAD settings, IPv6 addresses, and IPv6 prefixes. For more information, see the following sections:

- Configuring the Neighbor Solicitation Message Interval, page 7-19
- Configuring the Neighbor Reachable Time, page 7-19
- Configuring DAD Settings, page 7-20
- Configuring IPv6 Addresses on an Interface, page 7-20
- Configuring IPv6 Prefixes on an Interface, page 7-21

Configuring the Neighbor Solicitation Message Interval

You can configure the interval between IPv6 neighbor solicitation retransmissions on an interface. Valid values range from 1000 to 3600000 milliseconds. The default value is 1000 milliseconds. This setting is also sent in router advertisement messages.

To configure the neighbor solicitation message interval, perform the following steps:

Step 1 Choose **Configuration** > **Device Setup** > **Interfaces**.

Step 2 Choose the interface on which to configure the neighbor solicitation interval. The interface must have been configured with an IPv6 address. See the "Configuring IPv6 Neighbor Discovery" section on page 7-17 for more information.

Step 3 Click **Edit**. The Edit Interface dialog box appears with three tabs: General, Advanced, and IPv6.

Step 4 Click the **IPv6** tab.

Step 5 In the NS Interval field, enter the time interval.

Step 6 Click **OK**.

Step 7 Click **Apply** to save the configuration.

Configuring the Neighbor Reachable Time

The neighbor reachable time enables detecting unavailable neighbors. Shorter configured times enable detecting unavailable neighbors more quickly; however, shorter times consume more IPv6 network bandwidth and processing resources in all IPv6 network devices. Very short configured times are not recommended in normal IPv6 operation.

Valid time values range from 0 to 3600000 milliseconds. The default is 0; however, when you use 0, the reachable time is sent as undetermined. It is up to the receiving devices to set and track the reachable time value.

To configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event has occurred, perform the following steps:

Step 1 Choose **Configuration** > **Device Setup** > **Interfaces**.

Step 2 Choose the interface for which you want to configure the time. The interface must have been configured with an IPv6 address. For more information, see the "Configuring IPv6 Neighbor Discovery" section on page 7-17.

Step 3 Click **Edit**. The Edit Interface dialog box appears with three tabs: General, Advanced, and IPv6.

Step 4 Click the **IPv6** tab.

Step 5 In the Reachable Time field, enter a valid value.

Step 6 Click **OK**.

Step 7 Click **Apply** to save the configuration.

Configuring DAD Settings

Duplicate Address Detection (DAD) settings are part of the Neighbor Discovery configuration. DAD verifies the uniqueness of new unicast IPv6 addresses before they are assigned and ensures that duplicate IPv6 addresses are detected in the network on a link basis.

To specify DAD settings on the interface, perform the following steps:

Step 1 Enter the number of allowed DAD attempts. This setting configures the number of consecutive neighbor solicitation messages that are sent on an interface while DAD is performed on IPv6 addresses. Valid values are from 0 to 600. A zero value disables DAD processing on the specified interface. The default is one message.

Step 2 Enter the neighbor solicitation message interval. The neighbor solicitation message requests the link-layer address of a target node. Valid values are from 1000 to 3600000 milliseconds. The default is 1000 milliseconds.

Step 3 Enter the amount of time in seconds that a remote IPv6 node is considered reachable after a reachability confirmation event has occurred. Valid values are from 1000 to 3600000 milliseconds. The default is zero. A configured time enables the detection of unavailable neighbors. Shorter times enable detection more quickly; however, very short configured times are not recommended in normal IPv6 operation.

Step 4 Enter the amount of time that IPv6 router advertisement transmissions are considered valid. Valid values are from 3 to 1800 seconds. The default is 200 seconds. Router advertisement transmissions include a preference level and a lifetime field for each advertised router address. These transmissions provide route information and indicate that the router is still operational to network hosts. By default, these transmissions are sent every 400 to 600 seconds.

Step 5 Enter the interval between IPv6 router advertisement transmissions. Valid values are from 3 to 1800 seconds. The default is 200 seconds. To have the router advertisement transmission interval be listed in milliseconds, check the **RA Interval in Milliseconds** check box.

Step 6 To allow the generation of addresses for hosts, make sure that the Suppress RA check box is unchecked. This is the default setting if IPv6 unicast routing is enabled. To prevent the generation of IPv6 router advertisement transmissions, check the **Suppress RA** check box.

Step 7 To continue, see the "Configuring IPv6 Addresses on an Interface" section on page 7-20.

Configuring IPv6 Addresses on an Interface

To configure IPv6 addresses on an interface, perform the following steps:

Step 1 If you have not configured any IPv6 addresses with the CLI, to enable IPv6 addressing, check the **Enable IPv6** check box.

Step 2 To make sure that the source addresses of IPv6 packets received on that interface are verified according to the source MAC addresses to ensure that the interface identifiers use the modified EUI-64 format, check the **Enforce EUI-64** check box. If the interface identifiers do not conform to the modified EUI-64 format, an error message appears.

Step 3 If you are not going to assign any other IPv6 addresses, to set the link-local address manually, enter an address in the Link-local address field. A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. Alternatively, click the ellipsis to choose a link-local address from the Browse Link-local address dialog box.

Step 4 After you have selected the link-local address, click **OK** to return to the IPv6 tab.

The selected link-local address appears in the Link-local address field.

Step 5 To enable address autoconfiguration, check the **Enable address autoconfiguration** check box. During the stateless autoconfiguration process, duplicate address detection (DAD) verifies the uniqueness of new unicast IPv6 addresses before the addresses are assigned to interfaces (the new addresses remain in a tentative state while duplicate address detection is performed). Duplicate address detection is performed first on the new link-local address. When the link local address is verified as unique, then duplicate address detection is performed all the other IPv6 unicast addresses on the interface. For more information about DAD, see the "Configuring DAD Settings" section on page 7-20.

Step 6 In the Interface IPv6 Addresses area, click **Add**.

The Add IPv6 Address for Interface dialog box appears.

Step 7 (Optional) Check the **EUI-64** check box.

Step 8 Click **OK** to save your settings.

The Interface IPv6 Addresses Address field appears with the modified EUI-64 address.

Note You cannot use IPv6 addresses for the failover LAN and state links. For more information, see the "Configuring Failover with the High Availability and Scalability Wizard" section on page 17-4.

Step 9 To continue, see the "Configuring IPv6 Prefixes on an Interface" section on page 7-21.

Configuring IPv6 Prefixes on an Interface

To configure IPv6 prefixes on an interface, perform the following steps:

Step 1 In the Interface IPv6 Prefixes area, click **Add**.

The Add IPv6 Prefix for Interface dialog box appears.

Step 2 Enter the IPv6 address with the prefix length.

Step 3 (Optional) To configure the IPv6 address manually, check the **No Auto-Configuration** check box. This setting indicates to hosts on the local link that the specified prefix cannot be used for IPv6 autoconfiguration.

Step 4 (Optional) To indicate that the IPv6 prefix is not advertised, check the **No Advertisements** check box.

Step 5 (Optional) The **Off Link** check box indicates that the specified prefix is assigned to the link. Nodes sending traffic to addresses that contain the specified prefix consider the destination to be locally reachable on the link. This prefix should not be used for on-link determination.

Step 6 In the Prefix Lifetime area, click the **Lifetime Duration** radio button, and specify the following:

a. A valid lifetime for the prefix in seconds from the drop-down list. This setting is the amount of time that the specified IPv6 prefix is advertised as being valid. The maximum value represents infinity. Valid values are from 0 to 4294967295. The default is 2592000 (30 days).

b. A preferred lifetime for the prefix from the drop-down list. This setting is the amount of time that the specified IPv6 prefix is advertised as being preferred. The maximum value represents infinity. Valid values are from 0 to 4294967295. The default setting is 604800 (seven days).

Step 7 To define a prefix lifetime expiration date, click the **Lifetime Expiration Date** radio button, and specify the following:

a. Choose a valid month and day from the drop-down list, and then enter a time in hh:mm format.

b. Choose a preferred month and day from the drop-down list, and then enter a time in hh:mm format.

Step 8 Click **OK** to save your settings.

The Interface IPv6 Prefixes Address field appears with the preferred and valid dates.

Configuring Router Advertisement Messages

Router advertisement messages (ICMPv6 Type 134) are periodically sent from each IPv6 configured interface of the security appliance. The router advertisement messages are sent to the all-nodes multicast address Figure 7-2 shows an example of a router advertisement message.

Figure 7-2 IPv6 Neighbor Discovery—Router Advertisement Message

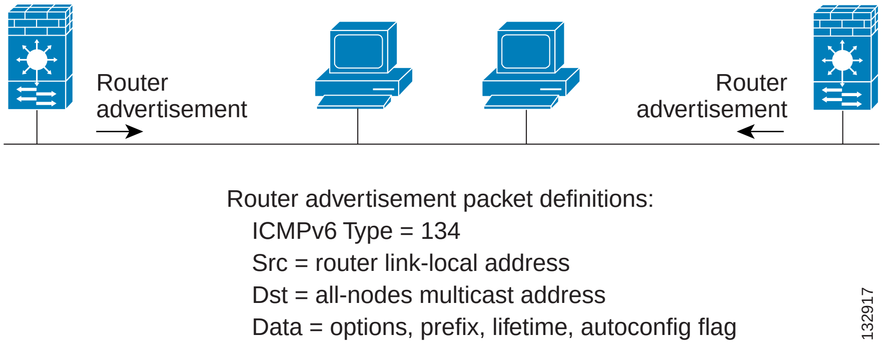
Router advertisement
Router advertisement
Router advertisement packet definitions:
ICMPv6 Type = 134
Src = router link-local address
Dst = all-nodes multicast address
Data = options, prefix, lifetime, autoconfig flag
132917

Router advertisement messages typically include the following information:

- One or more IPv6 prefix that nodes on the local link can use to automatically configure their IPv6 addresses.
- Lifetime information for each prefix included in the advertisement.
- Sets of flags that indicate the type of autoconfiguration (stateless or stateful) that can be completed.
- Default router information (whether the router sending the advertisement should be used as a default router and, if so, the amount of time (in seconds) the router should be used as a default router).
- Additional information for hosts, such as the hop limit and MTU a host should use in packets that it originates.
- The amount of time between neighbor solicitation message retransmissions on a given link.
- The amount of time a node considers a neighbor reachable.

Router advertisements are also sent in response to router solicitation messages (ICMPv6 Type 133). Router solicitation messages are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled router advertisement message. Because router solicitation messages are usually sent by hosts at system startup, and the host does not have a configured unicast address, the source address in router solicitation messages is usually the unspecified IPv6 address (0:0:0:0:0:0:0:0). If the host has a configured unicast address, the unicast address of the interface sending the router solicitation message is used as the source address in the message. The destination address in router solicitation messages is the all-routers multicast address with scope of the link. When a router advertisement is sent in response to a router solicitation, the destination address in the router advertisement message is the unicast address of the source of the router solicitation message.

You can configure the following settings for router advertisement messages:

- The time interval between periodic router advertisement messages.
- The router lifetime value, which indicates the amount of time IPv6 nodes should consider the security appliance to be the default router.
- The IPv6 network prefixes used on the link.
- Whether or not an interface transmits router advertisement messages.

Unless otherwise noted, the router advertisement message settings are specific to an interface and are entered in interface configuration mode. For information about changing these settings, see the following sections:

- Configuring the Router Advertisement Transmission Interval, page 7-23
- Configuring the Router Lifetime Value, page 7-24
- Suppressing Router Advertisement Messages, page 7-24

Configuring the Router Advertisement Transmission Interval

By default, router advertisements are sent out every 200 seconds. Valid values range from 3 to 1800 seconds.

The interval between transmissions should be less than or equal to the IPv6 router advertisement lifetime if the security appliance is configured as a default router. For more information, see the "Configuring the Router Lifetime Value" section on page 7-24. To prevent synchronization with other IPv6 nodes, randomly adjust the actual value used to within 20 percent of the desired value.

To change the interval between router advertisement transmissions on an interface, perform the following steps:

Step 1 Choose **Configuration** > **Device Setup** > **Interfaces**.

Step 2 Select the interface for which you want to configure the time.

The interface must have been configured with an IPv6 address. For more information, see the "Configuring IPv6 Neighbor Discovery" section on page 14-2.

Step 3 Click **Edit**. The Edit Interface dialog box appears with three tabs: General, Advanced, and IPv6.

Step 4 Click the **IPv6** tab.

Step 5 In the RA Interval field, enter a valid transmission interval value.

Note (Optional) To add a router advertisement transmission interval value in milliseconds instead, check the **RA Interval in Milliseconds** check box, and enter a value from 500 to 1800000.

Step 6 Click **OK**.

Step 7 Click **Apply** to save the configuration.

Configuring the Router Lifetime Value

The router lifetime value specifies how long nodes on the local link should consider the security appliance as the default router on the link. Valid values range from 0 to 9000 seconds. The default is 1800 seconds. Entering 0 indicates that the security appliance should not be considered a default router on the selected interface.

To configure the router lifetime value in IPv6 router advertisements on an interface, perform the following steps:

Step 1 Choose **Configuration** > **Device Setup** > **Interfaces**.

Step 2 Select the interface for which you want to configure the lifetime value.

The interface must have been configured with an IPv6 address. For more information see the "Configuring IPv6 Neighbor Discovery" section on page 14-2.

Step 3 Click **Edit**.

The Edit Interface dialog box appears with three tabs: General, Advanced, and IPv6.

Step 4 Click the **IPv6** tab.

Step 5 In the RA Lifetime field, enter a valid lifetime value.

Step 6 Click **OK**.

Step 7 Click **Apply** to save the configuration.

Suppressing Router Advertisement Messages

By default, router advertisement messages are automatically sent in response to router solicitation messages. You may want to disable these messages on any interface for which you do not want the security appliance to supply the IPv6 prefix (for example, the outside interface).

To suppress IPv6 router advertisement transmissions on an interface, perform the following steps:

Step 1 Choose **Configuration** > **Device Setup** > **Interfaces**.

Step 2 Select the interface for which you want to configure the lifetime value. The interface must have been configured with an IPv6 address. For more information, see the "Configuring IPv6 Neighbor Discovery" section on page 7-17.

Step 3 Click **Edit**.

The Edit Interface dialog box appears with three tabs: General, Advanced, and IPv6.

Step 4 Click the **IPv6** tab.

Step 5 Check the **Suppress RA** check box.

Step 6 Verify that the router advertisement message is suppressed on the interface that is configured for the IPv6 address.

Configuring IPv6 Static Neighbors

This section includes the following topics:

- Adding an IPv6 Static Neighbor, page 7-25
- Editing Static Neighbors, page 7-25
- Deleting Static Neighbors, page 7-26
- Viewing and Clearing Dynamic Neighbors, page 7-26

Adding an IPv6 Static Neighbor

Make sure that IPv6 is enabled on at least one interface before trying to add a neighbor, or ASDM returns an error message indicating that the configuration failed. For information about configuring IPv6 on an interface, see the "Configuring IPv6 Neighbor Discovery" section on page 14-2.

For information about configuring IPv6 Neighbor Discovery, see the "Configuring IPv6 Neighbor Discovery" section on page 7-17.

To add an IPv6 static neighbor, perform the following steps:

Step 1 Choose **Configuration** > **Device Management** > **Advanced** > **IPv6 Neighbor Discovery Cache**.

Step 2 Click **Add**.

The Add IPv6 Static Neighbor dialog box appears.

Step 3 From the Interface Name drop-down list, choose an interface on which to add the neighbor.

Step 4 In the IP Address field, enter the IPv6 address that corresponds to the local data-link address, or click the ellipsis (...) to browse for an address.

If an entry for the specified IPv6 address already exists in the neighbor discovery cache—learned through the IPv6 neighbor discovery process—the entry is automatically converted to a static entry.

Step 5 In the MAC address field, enter the local data-line (hardware) MAC address.

Step 6 Click **OK**.

Note Before you apply the changes and save the configuration, you can click **Reset** to cancel any changes and restore the original values.

Step 7 Click **Apply** to save the configuration.

Editing Static Neighbors

To edit a static neighbor that is defined in your configuration, perform the following steps:

Step 1 Choose **Configuration** > **Device Management** > **Advanced** > **IPv6 Neighbor Discovery Cache**.

Step 2 Select the neighbor from the main pane, and click **Edit**.

The Edit IPv6 Static Neighbor dialog box appears.

Step 3 Enter all necessary changes, and click **OK**.

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

Deleting Static Neighbors

To delete a static neighbor from your configuration, perform the following steps:

Step 1 Choose **Configuration** > **Device Management** > **Advanced** > **IPv6 Neighbor Discovery Cache**.

Step 2 Select the neighbor to delete from the main pane, and click **Delete**.

The selected neighbor is removed from the list.

Step 3 Click **Apply** to save the change to your current configuration.

Note Before you apply the changes and permanently delete the neighbor from your configuration, you can click **Reset** to restore the original values.

Viewing and Clearing Dynamic Neighbors

When a host or node communicates with a neighbor, the neighbor is added to the neighbor discovery cache. The neighbor is removed from the cache when there is no longer any communication with the neighbor.

To view dynamically discovered neighbors and to clear neighbors from the IPv6 Neighbor Discovery Cache, perform the following steps:

Step 1 Choose **Monitoring** > **Interface Graphs** > **IPv6 Neighbor Discovery Cache**.

You can view all static and dynamically discovered neighbors from the IPv6 Neighbor Discovery Cache pane.

Step 2 To clear all dynamically discovered neighbors from the cache, click **Clear Dynamic Neighbor Entries**.

The neighbor information is removed from the cache.

Note This procedure clears only dynamically discovered neighbors from the cache; it does not clear static neighbors. To clear static neighbors, see the "Deleting Static Neighbors" section on page 7-26.

Interface Configuration

To configure the remaining interfaces and enable traffic between two or more interfaces, perform the following steps:

Step 1 To change the configuration of the interface in the Edit Interface dialog box, click **Edit**.

Step 2 To enable traffic between two or more interfaces with the same security level, check the **Enable traffic between two or more interfaces with the same security level** check box.

Note IP address-related fields are not available in transparent mode.

Step 3 Click **Next** to continue.

Outside Interface Configuration - PPPoE

To configure the outside interface by obtaining an IP address from a PPPoE server, perform the following steps:

Step 1 Enter the name of the group. You must specify a group name to proceed.

Step 2 In the User Authentication area, enter the following information:

- The PPPoE username.
- The PPPoE password.
- Confirm the PPPoE password.

Step 3 In the Authentication Method area, enter the following:

- PAP is the default authentication method for PPPoE. You have the option of configuring CHAP or MS-CHAP manually. The username and password are sent unencrypted using this method.
- To select CHAP authentication, check the **CHAP** check box. CHAP does not prevent unauthorized access; it identifies the remote end. The access server then determines whether the user is allowed access.
- To select MS-CHAP authentication for PPP connections between a computer using a Windows operating system and an access server, check the **MS-CHAP** check box.

Step 4 In the IP Address area, choose one of the following:

- To obtain an IP address using a PPPoE server, click the **Obtain IP Address using PPPoE** radio button.
- To specify an IP address for an interface, click the **Specify an IP address** radio button.
  - Enter an IP address for an interface.
  - Enter or choose a subnet mask for an interface from the drop-down list.
- To obtain the default route between the PPPoE server and the PPPoE client, click the **Obtain default route using PPPoE** radio button.

Step 5 Click **Next** to continue.

Outside Interface Configuration

Note For all ASA 5500 series models except ASA 5505, with a full license, the security appliance supports up to five interfaces, with a maximum of three outside interfaces. In restricted mode, the security appliance supports up to three interfaces, and in transparent mode, the security appliance supports up to two interfaces. After you have created the maximum number of interfaces, or the maximum number of interfaces has already been named, you may not be able to create a new VLAN, and must select an existing one.

To configure the outside interface by specifying an IP address, or by obtaining one from a PPPoE or a DHCP server, perform the following steps:

Step 1 On the Interface Settings tab, do the following:

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

b. Add a name to a new interface or show the name associated with an existing interface.

c. To activate the interface in privileged mode, check the **Enable interface** check box.

d. Specify 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.

Step 2 Choose one of the following:

- To obtain an IP address from a PPPoE server, click the **Use PPPoE** radio button.
- To obtain an IP address from a DHCP server, click the **Use DHCP** radio button.
  - To obtain an IP address for the default gateway using DHCP, check the **Obtain default route using DHCP** check box.

Step 3 On the IPv6 Interface Settings tab, do the following:

- To enable the IPv6 interface, check the **Enable IPv6 for the Interface** check box.
- To add an IPv6 interface address, check the **Add an IPv6 Address** check box.
- Specify the IPv6 address and prefix length (for example, fe80:aabb::). Multiple addresses with prefixes are allowed; however, no two addresses can be repeated within the list of addresses.
  - To configure an address using the EUI 64-bit interface identifier format, check the **EUI 64** check box.

Step 4 Click **Next** to continue.

Feature History for the Startup Wizard

Table 7-3 lists the release history for this feature.

Table 7-3 Feature History for the Startup Wizard

Feature Name	Releases	Feature Information
Startup Wizard	ASA 7.0(1), ASDM 5.0(1)	This feature was introduced.

CHAPTER 8

Configuring Basic Device Settings

This section contains the following topics:

- Configuring the Management IP Address for Transparent Firewall Mode, page 8-1
- Setting the System Time, page 8-5
- Configuring Advanced Device Management Features, page 8-8
- Configuring the System Image/Configuration Settings, page 8-10
- Configuring the Device Name and Password, page 8-15
- Configuring System Software, page 8-16

Configuring the Management IP Address for Transparent Firewall Mode

This section describes how to configure the management IP address for transparent firewall mode, and includes the following topics:

- Information About the Management IP Address, page 8-1
- Licensing Requirements for the Management IP Address for a Transparent Firewall, page 8-2
- Guidelines and Limitations, page 8-2
- Configuring the IPv4 Address, page 8-3
- Configuring the IPv6 Address, page 8-3
- Configuring Advanced Device Management Features, page 8-8
- Feature History for the Management IP Address for a Transparent Firewall, page 8-5

Information About the Management IP Address

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. 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.

For IPv4 traffic, the management IP address is required to pass any traffic. For IPv6 traffic, you must, at a minimum, configure the link-local addresses to pass traffic, but a global management address is recommended for full functionality, including remote management and other management operations.

Note In addition to the management IP address for the device, 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. See Chapter 9, "Configuring Interfaces."

Although you do not configure IPv4 or global IPv6 addresses for other interfaces, you still need to configure the security level and interface name.

Licensing Requirements for the Management IP Address for a Transparent Firewall

Model	License Requirement
All models	Base License.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode. For multiple context mode, set the management IP address within each context.

Firewall Mode Guidelines

Supported in transparent firewall mode.

IPv6 Guidelines

- Supports IPv6.
- The following IPv6 address-related commands are not supported in transparent mode, because they require router capabilities:
  - **ipv6 address autoconfig**
  - **ipv6 nd suppress-ra**
- No support for IPv6 anycast addresses.
- You can configure both IPv6 and IPv4 addresses.

Additional Guidelines and Limitations

- In addition to the management IP address for the device, 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. See Chapter 9, "Configuring Interfaces."

- Although you do not configure IP addresses for other interfaces, you still need to configure the security level and interface name according to Chapter 9, "Configuring Interfaces."

Configuring the IPv4 Address

To set the management IPv4 address, perform the following steps:

Step 1 Go to Configuration > Device Management > Management Access > Management IP Address.

Step 2 In the IPv4 Address area, enter the IP address in the Management IP Address field.

This address must be on the same subnet as the upstream and downstream routers. You cannot set the subnet to a host subnet (255.255.255.255). The **standby** keyword and address is used for failover. See the "Failover: Interfaces (Transparent Firewall Mode)" section on page 17-19 for more information.

Step 3 From the Subnet Mask drop-down list, choose a subnet mask, or enter a subnet mask directly in the field.

Step 4 Click **Apply**.

Configuring the IPv6 Address

You can configure two types of unicast addresses for IPv6:

- Global—The global address is a public address that you can use on the public network. This address needs to be configured for the whole device, and not per-interface.
- Link-local—The link-local address is a private address that you can only use on the directly-connected network. Routers do not forward packets using link-local addresses; they are only for communication on a particular physical network segment. They can be used for address configuration or for the ND functions such as address resolution and neighbor discovery. Because the link-local address is only available on a segment, and is tied to the interface MAC address, you need to configure the link-local address per interface.

At a minimum, you need to configure a link-local addresses for IPv6 to operate. If you configure a global address, a link-local addresses is automatically configured on each interface, so you do not also need to specifically configure a link-local address. If you do not configure a global address, then you need to configure the link-local address, either automatically or manually.

See the "IPv6 Addresses" section on page B-5 for more information about IPv6 addressing.

This section describes how to configure the global address or the link-local address, and includes the following topics:

- Configuring the Link-Local Addresses Automatically, page 8-4
- Configuring the Link-Local Addresses Automatically, page 8-4
- Configuring DAD Settings, page 8-4

Configuring the Global Address

To set the management IPv6 address, perform the following steps:

Step 1 Go to Configuration > Device Management > Management Access > Management IP Address.

Step 2 In the IPv6 Addresses area, click **Add**.

The Add IPv6 Management Address dialog box appears.

Step 3 In the IP Address field, enter an IPv6 address.

For example, 2001:0DB8::BA98:0:3210. See the "IPv6 Addresses" section on page B-5 for more information about IPv6 addressing.

Step 4 In the Prefix Length field, enter the prefix length.

For example, 48. See the "IPv6 Addresses" section on page B-5 for more information about IPv6 addressing.

Step 5 Click **OK**.

Step 6 To configure additional addresses, repeat Step 2 through Step 5.

Step 7 Click **Apply**.

Configuring the Link-Local Addresses Automatically

If you only need to configure a link-local address and are not going to assign any other IPv6 addresses, you have the option of generating the link-local addresses based on the interface MAC addresses (Modified EUI-64 format). To manually assign the link-local address, see the "Configuring the Link-Local Address on an Interface (Transparent Firewall Mode)" section on page 9-30.

To automatically configure the link-local addresses for all interfaces, perform the following steps:

Step 1 Go to Configuration > Device Management > Management Access > Management IP Address.

Step 2 In the IPv6 configuration area, check **Enable IPv6**.

This option enables IPv6 on all interfaces and automatically generates the link-local addresses using the Modified EUI-64 interface ID based on the interface MAC address.

Note You do not need to check this option if you configure any IPv6 addresses (either global or link-local); IPv6 support is automatically enabled as soon as you assign an IPv6 address. Similarly, unchecking this option does not disable IPv6 if you configured IPv6 addresses.

To configure IPv6 DAD parameters, shown in this area, see the "Configuring DAD Settings" section on page 8-4.

Step 3 Click **Apply**.

Configuring DAD Settings

DAD verifies the uniqueness of new unicast IPv6 addresses before they are assigned and ensures that duplicate IPv6 addresses are detected in the network on a link basis.

For information about the Enable IPv6 parameter, see the "Configuring the Link-Local Addresses Automatically" section on page 8-4.

To specify DAD settings on an interface, perform the following steps:

Step 1 Go to Configuration > Device Management > Management Access > Management IP Address.

Step 2 In the IPv6 configuration area, in the DAD attempts field, enter the number of allowed DAD attempts.

This setting configures the number of consecutive neighbor solicitation messages that are sent on an interface while DAD is performed on IPv6 addresses. Valid values are from 0 to 600. A zero value disables DAD processing on the specified interface. The default is one message.

Step 3 In the NS Interval field, enter the neighbor solicitation message interval.

The neighbor solicitation message requests the link-layer address of a target node. Valid values are from 1000 to 3600000 milliseconds. The default is 1000 milliseconds.

Step 4 In the Reachable Time field, enter the amount of time in seconds that a remote IPv6 node is considered reachable after a leachability confirmation event has occurred.

Valid values are from 1000 to 3600000 milliseconds. The default is zero. A configured time enables the detection of unavailable neighbors. Shorter times enable detection more quickly; however, very short configured times are not recommended in normal IPv6 operation.

Step 5 Click **Apply**.

Feature History for the Management IP Address for a Transparent Firewall

Table 8-1 lists the release history for this feature.

Table 8-1 Feature History for Transparent Mode Management Address

Feature Name	Releases	Feature Information
IPv6 support	8.2(1)	IPv6 support was introduced for transparent firewall mode.

Setting the System Time

You can manually set the system date or time or have the security appliance dynamically set the system date and tim e using an NTP server.

See the following topics for more information:

- Clock, page 8-5
- NTP, page 8-6

Clock

The Clock pane 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 pane.

In multiple context mode, you can set the time in the system configuration only.

To dynamically set the time using an NTP server, see the **NTP** pane; time derived from an NTP server overrides any time set manually in the **Clock** pane.

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 savings time, from 2:00 a.m. on the second Sunday in March to 2:00 a.m. on the first Sunday in November.

Note Changing the time zone on the security appliance may drop the connection to intelligent SSMs.

- Date—Sets the date. Click the Date drop-down list to display a calendar. Then navigate to the correct date using the following methods:
  - Click the name of the month to display a list of months. Click the desired month. The calendar updates to that month.
  - Click the year to change the year. You can use the up and down arrows to scroll through the years, or you can type a year in the entry field.
  - Click the arrows to the right and left of the month and year display to scroll the calendar forward and backwards one month at a time.
  - Click a day on the calendar to set the date.
- 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 in the bottom right corner of the ASDM pane. 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
•	•	•	—	•

NTP

The NTP pane 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 pane.

Time derived from an NTP server overrides any time set manually in the **Clock** pane.

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

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
•	•	•	—	•

Configuring Advanced Device Management Features

The following sections describe how to configure the items in the Advanced section.

- Configuring HTTP Redirect, page 8-8
- History Metrics, page 8-9
- Configuring IPv6 Neighbor Discovery, page 14-2
- Creating the Standard Priority Queue for an Interface, page 26-5
- Configuring SSL Settings, page 41-1
- WCCP, page 15-13

Configuring HTTP Redirect

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

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

To change the HTTP redirect setting of an interface or the port from which it redirects HTTP connections, select the interface in the table and click **Edit**. You can also double-click an interface. The Edit HTTP/HTTPS Settings dialog box opens.

Edit HTTP/HTTPS Settings

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

Fields

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

- Interface—Identifies the interface on which the security appliance redirects or does not redirect HTTP requests to HTTPS.
- Redirect HTTP to HTTPS—Check to redirect HTTP requests to HTTPS, or uncheck to not redirect HTTP requests to HTTPS.
- HTTP Port—Identifies the port from which the interface redirects HTTP connections. By default it listens to port 80.

Modes

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

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

Configuring Maximum SSL VPN Sessions

This pane lets you set a maximum number of SSL VPN sessions.

Fields

Maximum Sessions—Enter the maximum number of Clientless SSL VPN sessions you want to allow. Be aware that the different ASA models support Clientless SSL VPN sessions as follows: ASA 5510 supports a maximum of 250; ASA 5520 maximum is 750; ASA 5540 maximum is 2500; ASA 5550 maximum is 5000.

Modes

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

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

History Metrics

The History Metrics pane lets you configure the adaptive security appliance to keep a history of various statistics, which ASDM can display 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.

To configure history metrics, perform the following steps:

Step 1 Choose **Configuration** > **Device Management** > **Advanced** > **History Metrics**.

The History Metrics pane appears.

Step 2 Check the **ASDM History Metrics** check box to enable history metrics, and then click **Apply**.

Modes

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

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

Configuring the System Image/Configuration Settings

This section includes the following topics:

- Configuring Auto Update, page 8-10
- Configuring the Boot Image/Configuration Settings, page 8-13

Configuring Auto Update

The Auto Update pane lets you configure the security appliance to be managed remotely from servers that support 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.

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 order 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 box, 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 time out. The Timeout area contains the following fields:

- Enable Timeout Period—Check to enable the security appliance to time out if no response is received from the Auto Update server.
- Timeout Period (Minutes)—Enter the number of minutes the security appliance will wait to time out 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 box 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

The Set Polling Schedule dialog box 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 box 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 pane 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 randomization—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

The Edit Auto Update Server dialog box 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

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
•	•	•	—	—

Configuring the Boot Image/Configuration Settings

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 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 Configuration 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

To add a boot image entry to the boot order list, click **Add** in the Boot Image/Configuration pane.

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 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
•	•	•	—	•

Configuring the Device Name and Password

The Device Name/Password pane lets you set the hostname and domain name for the security appliance and set the enable and telnet passwords.

The hostname appears in 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; it 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.”

The Telnet Password sets the login password. By default, it is “cisco.” Although this area is called Telnet Password, this password applies to Telnet and SSH access. 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.)

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 pane. (If you configure user authentication for enable access, then each user has their own password, and this enable password is not used; see About Authentication, page 16-2. In addition, you can configure authentication for HTTP/ASDM access.)

Fields

The Hostname and Domain Name area contains the following fields:

- Hostname—Sets the hostname. The default hostname depends on your platform.
- Domain Name—Sets the domain name. The default domain name is default.domain.invalid.

The Enable Password area contains the following fields. In multiple context mode, the Enable Password area only appears in contexts; it does not appear in the system execution space.

- 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.

The Telnet Password area contains the following fields. In multiple context mode, the Telnet Password area only appears in contexts; it does not appear in the system execution space.

- 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
•	•	•	•	•

Configuring System Software

The System Software pane lets you configure the parameters of security appliances configured as Auto Update clients when this security appliance is acting 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.

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 box, 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

Fields

The Add/Edit Client Update dialog box 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 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
•	•	•	—	—

CHAPTER 9

Configuring Interfaces

This chapter describes how to configure interfaces, including Ethernet parameters, switch ports (for the ASA 5505), VLAN subinterfaces, and IP addressing.

The procedure to configure interfaces varies depending on several factors: the ASA 5505 vs. other models; routed vs. transparent mode; and single vs. multiple mode. This chapter describes how to configure interfaces for each of these variables.

Note If your security appliance has the default factory configuration, many interface parameters are already configured. This chapter assumes you do *not* have a factory default configuration, or that if you have a default configuration, that you need to change the configuration. For information about the factory default configurations, see the "Factory Default Configurations" section on page 1-1.

This chapter includes the following sections:

- Information About Interfaces, page 9-1
- Licensing Requirements for Interfaces, page 9-6
- Guidelines and Limitations, page 9-6
- Default Settings, page 9-7
- Starting Interface Configuration (ASA 5510 and Higher), page 9-8
- Starting Interface Configuration (ASA 5505), page 9-15
- Completing Interface Configuration (All Models), page 9-20
- Allowing Same Security Level Communication, page 9-31
- Enabling Jumbo Frame Support (ASA 5580, Multiple Mode), page 9-32
- Monitoring Interfaces, page 9-32
- Feature History for Interfaces, page 9-33

Information About Interfaces

This section describes security appliance interfaces, and includes the following topics:

- ASA 5505 Interfaces, page 9-2
- Auto-MDI/MDIX Feature, page 9-4
- Security Levels, page 9-5

- Dual IP Stack, page 9-5
- Management Interface (ASA 5510 and Higher), page 9-6

ASA 5505 Interfaces

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

- Understanding ASA 5505 Ports and Interfaces, page 9-2
- Maximum Active VLAN Interfaces for Your License, page 9-2
- VLAN MAC Addresses, page 9-4
- Power Over Ethernet, page 9-4

Understanding ASA 5505 Ports and Interfaces

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

- Physical switch ports—The security appliance has 8 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 "Power Over Ethernet" section on page 9-4 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 security appliance applies the security policy to the traffic and routes or bridges between the two VLANs.

Maximum Active VLAN Interfaces for Your License

In transparent firewall mode, you can configure the following VLANs depending on your license:

- Base license—2 active VLANs.
- Security Plus license—3 active VLANs, one of which must be for failover.

In routed mode, you can configure the following VLANs depending on your license: Base license

- Base license—3 active VLANs. The third VLAN can only be configured to initiate traffic to one other VLAN. See Figure 9-1 for more information.
- Security Plus license—20 active VLANs.

Note 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 9-1 for an example network where the Home VLAN can communicate with the Internet, but cannot initiate contact with Business.

Figure 9-1 ASA 5505 Adaptive Security Appliance with Base License

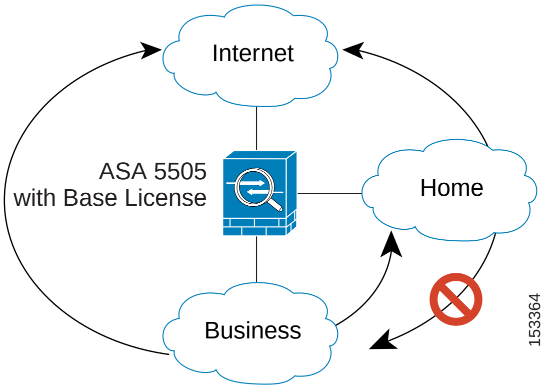

With the Security Plus license, you can configure 20 VLAN interfaces, including a VLAN interface for failover and a VLAN interface as a backup link to your ISP. You can configure the backup interface to not pass through traffic unless the route through the primary interface fails. You can configure trunk ports to accommodate multiple VLANs per port.

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

See Figure 9-2 for an example network.

Figure 9-2 ASA 5505 Adaptive Security Appliance with Security Plus License

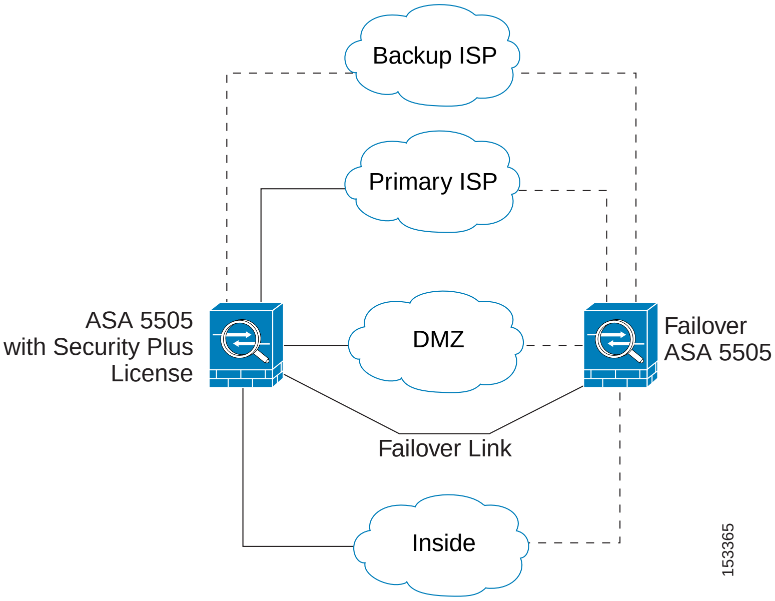

VLAN MAC Addresses

- 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. See the "Configuring Advanced Interface Parameters" section on page 9-25.
- Transparent firewall mode—Each VLAN has a unique MAC address. You can override the generated MAC addresses if desired by manually assigning MAC addresses. See the "Configuring Advanced Interface Parameters" section on page 9-25.

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 security appliance does not supply power to the switch ports.

If you shut down the switch port, you disable power to the device. Power is restored when you enable the portd. See the "Configuring and Enabling Switch Ports as Access Ports" section on page 9-17 for more information about shutting down a switch port.

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.

You can only enable SPAN monitoring using the Command Line Interface tool by entering the **switchport monitor** command. See the **switchport monitor** command in the *Cisco ASA 5500 Series Command Reference* for more information.

ASA 5580 Interfaces

The ASA 5580 security appliance supports multiple types of Ethernet interfaces including Gigabit Ethernet and 10-Gigabit Ethernet speeds, and copper and fiber connectors. See the *Cisco ASA 5580 Adaptive Security Appliance Getting Started Guide* for detailed information about the interface adapters available for the ASA 5580 security appliance, and which slots support each adapter type.

Auto-MDI/MDIX Feature

For RJ-45 interfaces on the ASA 5500 series 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. For Gigabit Ethernet, when the speed and duplex are set to 1000 and full, then the interface always auto-negotiates; therefore Auto-MDI/MDIX is always enabled and you cannot disable it.

Security Levels

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. See the "Allowing Same Security Level Communication" section on page 9-31 for more information.

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.

  If you enable communication for same security interfaces (see the "Allowing Same Security Level Communication" section on page 9-31), 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).

  If you enable communication 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.

  If you enable communication for same security interfaces, you can configure **established** commands for both directions.

Dual IP Stack

The security appliance supports the configuration of both IPv6 and IPv4 on an interface. You do not need to enter any special commands to do so; simply enter the IPv4 configuration commands and IPv6 configuration commands as you normally would. Make sure you configure a default route for both IPv4 and IPv6.

Management Interface (ASA 5510 and Higher)

The management interface is a Fast Ethernet interface designed for management traffic only. You can, however, use it for through traffic if desired. In transparent firewall mode, you can use the management interface (for management purposes) in addition to the two interfaces allowed for through traffic. You can also add subinterfaces to the management interface to provide management in each security context for multiple context mode.

Licensing Requirements for Interfaces

The following table shows the licensing requirements for VLANs:

Model	License Requirement
ASA 5505	Base License: 3 (2 regular zones and 1 restricted zone that can only communicate with 1 other zone) Security Plus License: 20
ASA 5510	Base License: 50 Security Plus License: 100
ASA 5520	Base License: 150
ASA 5540	Base License: 200
ASA 5550	Base License: 250
ASA 5580	Base License: 250

The following table shows the licensing requirements for VLAN trunks:

Model	License Requirement
ASA 5505	Base License: None. Security Plus License: 8.
All other models	N/A

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

In multiple context mode, configure the physical interfaces in the system execution space according to the "Starting Interface Configuration (ASA 5510 and Higher)" section on page 9-8.

Then, configure the logical interface parameters in the context execution space according to the "Completing Interface Configuration (All Models)" section on page 9-20.

Firewall Mode Guidelines

Transparent firewall mode allows only two interfaces to pass through traffic; however, on the ASA 5510 and higher security appliance, you can use the Management 0/0 or 0/1 interface (either the physical interface or a subinterface) as a third interface for management traffic. The mode is not configurable in this case and must always be management-only.

Failover Guidelines

Do not finish configuring failover interfaces with the procedures in "Completing Interface Configuration (All Models)" section on page 9-20. See the "Failover Link Configuration" section on page 17-10 and "State Link Configuration" section on page 17-11 to configure the failover and state links. In multiple context mode, failover interfaces are configured in the system configuration.

IPv6 Guidelines

Supports IPv6.

In transparent mode on a per interface basis, you can only configure the link-local address; you configure the global address as the management address for the entire unit, but not per interface. Because configuring the management global IP address automatically configures the link-local addresses per interface, the only IPv6 configuration you need to perform is to set the management IP address according to the "Configuring the IPv6 Address" section on page 8-3.

Model Guidelines

Subinterfaces are not available for the ASA 5505 security appliance.

Default Settings

This section lists default settings for interfaces if you do not have a factory default configuration. For information about the factory default configurations, see the "Factory Default Configurations" section on page 1-1.

Default Security Level

The default security level is 0. If you name an interface "inside" and you do not set the security level explicitly, then the security appliance sets the security level to 100.

Note If you change the security level of an interface, and you do not want to wait for existing connections to time out before the new security information is used, you can clear the connections using the **clear local-host** command.

Default State of Interfaces

The default state of an interface depends on the type and the context mode.

In multiple context mode, all allocated interfaces are enabled by default, no matter what the state of the interface is in the system execution space. However, for traffic to pass through the interface, the interface also has to be enabled in the system execution space. If you shut down an interface in the system execution space, then that interface is down in all contexts that share it.

In single mode or in the system execution space, interfaces have the following default states:

- Physical interfaces and switch ports—Disabled.

- Redundant Interfaces—Enabled. However, for traffic to pass through the redundant interface, the member physical interfaces must also be enabled.
- Subinterfaces or VLANs—Enabled. However, for traffic to pass through the subinterface, the physical interface must also be enabled.

Default Speed and Duplex

- By default, the speed and duplex for copper (RJ-45) interfaces are set to auto-negotiate.
- The fiber interface for the ASA 5550 and the 4GE SSM has a fixed speed and does not support duplex, but you can set the interface to negotiate link parameters (the default) or not to negotiate.
- For fiber interfaces for the ASA 5580, the speed is set for automatic link negotiation.

Default Connector Type

The ASA 5550 security appliance and the 4GE SSM for the ASA 5510 and higher security appliance include two connector types: copper RJ-45 and fiber SFP. RJ-45 is the default. You can configure the security appliance to use the fiber SFP connectors.

Default 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.

Starting Interface Configuration (ASA 5510 and Higher)

This section includes tasks for starting your interface configuration for the ASA 5510 and higher.

Note For multiple context mode, complete all tasks in this section in the system execution space. If you are not already in the system execution space , in the Configuration > Device List pane, double-click **System** under the active device IP address.

For ASA 5505 configuration, see the "Starting Interface Configuration (ASA 5505)" section on page 9-15.

This section includes the following topics:

Task Flow for Starting Interface Configuration

To start configuring interfaces, perform the following steps:

Step 1 (Multiple context mode) Complete all tasks in this section in the system execution space. If you are not already in the System configuration mode, in the Configuration > Device List pane, double-click **System** under the active device IP address.

Step 2 Enable the physical interface, and optionally change Ethernet parameters. See the "Enabling the Physical Interface and Configuring Ethernet Parameters" section on page 9-9.

Physical interfaces are disabled by default.

Step 3 (Optional) Configure redundant interface pairs. See the "Configuring a Redundant Interface" section on page 9-11.

A logical redundant interface pairs an active and a standby physical interface. When the active interface fails, the standby interface becomes active and starts passing traffic.

Step 4 (Optional) Configure VLAN subinterfaces. See the "Configuring VLAN Subinterfaces and 802.1Q Trunking" section on page 9-13.

Step 5 (Multiple context mode only) Assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See the"Assigning Interfaces to Contexts and Automatically Assigning MAC Addresses (Multiple Context Mode)" section on page 9-15.

Step 6 Complete the interface configuration according to the "Completing Interface Configuration (All Models)" section on page 9-20.

Enabling the Physical Interface and Configuring Ethernet Parameters

This section describes how to:

- Enable the physical interface
- Set a specific speed and duplex (if available)
- Enable pause frames for flow control (ASA 5580 10 Gigabit Ethernet only).

Prerequisites

For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the Configuration > Device List pane, double-click **System** under the active device IP address.

Detailed Steps

Step 1 Depending on your context mode:

- For single mode, choose the **Configuration** > **Device Setup** > **Interfaces** pane.
- For multiple mode in the System execution space, choose the **Configuration** > **Context Management** > **Interfaces** pane.

By default, all physical interfaces are listed.

Step 2 Click a physical interface that you want to configure, and click **Edit**.

The Edit Interface dialog box appears.

Step 3 To enable the interface, check the **Enable Interface** check box.

Step 4 To add a description, enter text in the Description field.

The description can be up to 240 characters on a single line, without carriage returns. 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.

Step 5 (Optional) To set the media type, duplex, and speed, or for the ASA 5580 10 Gigabit Ethernet interface, enable pause frames for flow control, click **Configure Hardware Properties**.

a. If you have an ASA 5550 adaptive security appliance or a 4GE SSM, you can choose either **RJ-45** or **SFP** from the Media Type drop-down list.

RJ-45 is the default.

b. To set the duplex for RJ-45 interfaces, choose **Full**, **Half**, or **Auto**, depending on the interface type, from the Duplex drop-down list.

c. To set the speed, choose a value from the Speed drop-down list.

The speeds available depend on the interface type. For SFP interfaces, 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. See the "Auto-MDI/MDIX Feature" section on page 9-4.

d. For the ASA 5580 10 Gigabit Ethernet interfaces, to enable pause (XOFF) frames for flow control on 10 Gigabit Ethernet interfaces, check the **Enable Pause Frame** check box.

If you have a traffic burst, dropped packets can occur if the burst exceeds the buffering capacity of the FIFO buffer on the NIC and the receive ring buffers. Enabling pause frames for flow control can alleviate this issue. Pause (XOFF) and XON frames are generated automatically by the NIC hardware based on the FIFO buffer usage. A pause frame is sent when the buffer usage exceeds the High Watermark. The default value is 128 KB; you can set it between 0 and 511. After a pause is sent, an XON frame can be sent when the buffer usage is reduced below the Low Watermark. By default, the value is 64 KB; you can set it between 0 and 511. The link partner can resume traffic after receiving an XON, or after the XOFF expires, as controlled by the Pause Time value in the pause frame. The default value is 26624; you can set it between 0 and 65535. If the buffer usage is consistently above the High Watermark, pause frames are sent repeatedly, controlled by the pause refresh threshold value.

To change the default values for the Low Watermark, High Watermark, and Pause Time, uncheck the **Use Default Values** check box.

Note Only flow control frames defined in 802.3x are supported. Priority-based flow control is not supported.

e. Click **OK** to accept the Hardware Properties changes.

Step 6 Click **OK** to accept the Interface changes.

What to Do Next

Optional Tasks:

- Configure redundant interface pairs. See the "Configuring a Redundant Interface" section on page 9-11.

- Configure VLAN subinterfaces. See the "Configuring VLAN Subinterfaces and 802.1Q Trunking" section on page 9-13.

Required Tasks:

- For multiple context mode, assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See the "Assigning Interfaces to Contexts and Automatically Assigning MAC Addresses (Multiple Context Mode)" section on page 9-15.
- For single context mode, complete the interface configuration. See the "Completing Interface Configuration (All Models)" section on page 9-20.

Configuring a Redundant Interface

A logical redundant interface consists of a pair of physical interfaces: an active and a standby interface. When the active interface fails, the standby interface becomes active and starts passing traffic. You can configure a redundant interface to increase the security appliance reliability. This feature is separate from device-level failover, but you can configure redundant interfaces as well as failover if desired.

This section describes how to configure redundant interfaces, and includes the following topics:

- Configuring a Redundant Interface, page 9-11
- Changing the Active Interface, page 9-13

Configuring a Redundant Interface

This section describes how to create a redundant interface. By default, redundant interfaces are enabled.

Guidelines and Limitations

- You can configure up to 8 redundant interface pairs.
- All security appliance configuration refers to the logical redundant interface instead of the member physical interfaces.
- Redundant interface delay values are configurable, but by default the security appliance will inherit the default delay values based on the physical type of its member interfaces.
- The only configuration available to physical interfaces that are part of a redundant interface pair are physical parameters (set in the "Enabling the Physical Interface and Configuring Ethernet Parameters" section on page 9-9).
- If you shut down the active interface, then the standby interface becomes active.

For failover, follow these guidelines when adding member interfaces:

- If you want to use a redundant interface for the failover or state link, then you must configure the redundant interface as part of the basic configuration on the secondary unit in addition to the primary unit.
- If you use a redundant interface for the failover or state link, you must put a switch or hub between the two units; you cannot connect them directly. Without the switch or hub, you could have the active port on the primary unit connected directly to the standby port on the secondary unit.
- You can monitor redundant interfaces for failover.
- When the active interface fails over to the standby interface, this activity does not cause the redundant interface to appear to be failed when being monitored for device-level failover. Only when both physical interfaces fail does the redundant interface appear to be failed.

Redundant Interface MAC Address

The redundant interface uses the MAC address of the first physical interface that you add. If you change the order of the member interfaces in the configuration, then the MAC address changes to match the MAC address of the interface that is now listed first. Alternatively, you can assign a MAC address to the redundant interface, which is used regardless of the member interface MAC addresses (see the "Configuring Advanced Interface Parameters" section on page 9-25 or the "Assigning Interfaces to Contexts and Automatically Assigning MAC Addresses (Multiple Context Mode)" section on page 9-15). When the active interface fails over to the standby, the same MAC address is maintained so that traffic is not disrupted.

Prerequisites

- Both member interfaces must be of the same physical type. For example, both must be Ethernet.
- You cannot add a physical interface to the redundant interface if you configured a name for it. You must first remove the name in the Configuration > Device Setup > Interfaces pane.
- For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the Configuration > Device List pane, double-click **System** under the active device IP address.

Caution If you are using a physical interface already in your configuration, removing the name will clear any configuration that refers to the interface.

Detailed Steps

You can configure up to 8 redundant interface pairs. To configure a redundant interface, perform the following steps:

Step 1 Depending on your context mode:

- For single mode, choose the **Configuration** > **Device Setup** > **Interfaces** pane.
- For multiple mode in the System execution space, choose the **Configuration** > **Context Management** > **Interfaces** pane.

Step 2 Click **Add** > **Redundant Interface**.

The Add Redundant Interface dialog box appears.

Step 3 In the Redundant ID field, enter an integer between 1 and 8.

Step 4 From the Primary Interface drop-down list, choose the physical interface you want to be primary.

Be sure to pick an interface that does not have a subinterface and that has not already been allocated to a context.

Step 5 From the Secondary Interface drop-down list, choose the physical interface you want to be secondary.

Step 6 If the interface is not already enabled, check the **Enable Interface** check box.

The interface is enabled by default. To disable it, uncheck the check box.

Step 7 To add a description, enter text in the Description field.

The description can be 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.

Step 8 Click **OK**.

You return to the Interfaces pane.

What to Do Next

Optional Task:

- Configure VLAN subinterfaces. See the "Configuring VLAN Subinterfaces and 802.1Q Trunking" section on page 9-13.

Required Tasks:

- For multiple context mode, assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See the "Assigning Interfaces to Contexts and Automatically Assigning MAC Addresses (Multiple Context Mode)" section on page 9-15.
- For single context mode, complete the interface configuration. See the "Completing Interface Configuration (All Models)" section on page 9-20.

Changing the Active Interface

By default, the active interface is the first interface listed in the configuration, if it is available. To view which interface is active, enter the following command in the Tools > Command Line Interface tool:

```
show interface redundantnumber detail | grep Member
```

For example:

```
show interface redundant1 detail | grep Member
      Members GigabitEthernet0/3(Active), GigabitEthernet0/2
```

To change the active interface, enter the following command:

```
redundant-interface redundantnumber active-member physical_interface
```

where the **redundant***number* argument is the redundant interface ID, such as **redundant1**.

The *physical_interface* is the member interface ID that you want to be active.

Configuring VLAN Subinterfaces and 802.1Q Trunking

Subinterfaces let you divide a physical or redundant interface into multiple logical interfaces that are tagged with different VLAN IDs. An interface with one or more VLAN subinterfaces is automatically configured as an 802.1Q trunk. Because VLANs allow you to keep traffic separate on a given physical interface, you can increase the number of interfaces available to your network without adding additional physical interfaces or security appliances. This feature is particularly useful in multiple context mode so that you can assign unique interfaces to each context.

Guidelines and Limitations

- Maximum subinterfaces—To determine how many VLAN subinterfaces are allowed for your platform, see the "Licensing Requirements for Interfaces" section on page 9-6.
- Preventing untagged packets on the physical interface—If you use subinterfaces, you typically do not also want the physical interface to pass traffic, because the physical interface passes untagged packets. This property is also true for the active physical interface in a redundant interface pair. Because the physical or redundant interface must be enabled for the subinterface to pass traffic, ensure that the physical or redundant interface does not pass traffic by not configuring a name for the interface. If you want to let the physical or redundant interface pass untagged packets, you can configure the name as usual. See the "Completing Interface Configuration (All Models)" section on page 9-20 for more information about completing the interface configuration.

Prerequisites

For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the Configuration > Device List pane, double-click **System** under the active device IP address.

Detailed Steps

To add a subinterface and assign a VLAN to it, perform the following steps:

Step 1 Depending on your context mode:

- For single mode, choose the **Configuration** > **Device Setup** > **Interfaces** pane.
- For multiple mode in the System execution space, choose the **Configuration** > **Context Management** > **Interfaces** pane.

Step 2 Click **Add** > **Interface**.

The Add Interface dialog box appears.

Step 3 From the Hardware Port drop-down list, choose the physical interface to which you want to add the subinterface.

Step 4 If the interface is not already enabled, check the **Enable Interface** check box.

The interface is enabled by default. To disable it, uncheck the check box.

Step 5 In the VLAN ID field, enter 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.

Step 6 In the Subinterface ID field, enter 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.

Step 7 (Optional) In the Description field, enter a description for this interface.

The description can be 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.

Step 8 Click **OK**.

You return to the Interfaces pane.

What to Do Next

- For multiple context mode, assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See the "Assigning Interfaces to Contexts and Automatically Assigning MAC Addresses (Multiple Context Mode)" section on page 9-15.
- For single context mode, complete the interface configuration. See the "Completing Interface Configuration (All Models)" section on page 9-20.

Assigning Interfaces to Contexts and Automatically Assigning MAC Addresses (Multiple Context Mode)

To complete the configuration of interfaces in the system execution space, perform the following tasks that are documented in Chapter 11, "Configuring Security Contexts":

- To assign interfaces to contexts, see the "Configuring Security Contexts" section on page 11-16.
- (Optional) To automatically assign unique MAC addresses to context interfaces, see the "Automatically Assigning MAC Addresses" section on page 11-18.

  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. Alternatively, you can manually assign MAC addresses within the context according to the "Configuring Advanced Interface Parameters" section on page 9-25.

What to Do Next

Complete the interface configuration. See the "Completing Interface Configuration (All Models)" section on page 9-20.

Starting Interface Configuration (ASA 5505)

This section includes tasks for starting your interface configuration for the ASA 5505 security appliance, including creating VLAN interfaces and assigning them to switch ports. See the "Understanding ASA 5505 Ports and Interfaces" section on page 9-2 for more information.

For ASA 5510 and higher configuration, see the "Starting Interface Configuration (ASA 5510 and Higher)" section on page 9-8.

This section includes the following topics:

- Task Flow for Starting Interface Configuration, page 9-16
- Configuring VLAN Interfaces, page 9-16
- Configuring and Enabling Switch Ports as Access Ports, page 9-17
- Configuring and Enabling Switch Ports as Trunk Ports, page 9-19

Task Flow for Starting Interface Configuration

To configure interfaces in single mode, perform the following steps:

Step 1 Configure VLAN interfaces. See the "Configuring VLAN Interfaces" section on page 9-16.

Step 2 Configure and enable switch ports as access ports. See the "Configuring and Enabling Switch Ports as Access Ports" section on page 9-17.

Step 3 (Optional for Security Plus licenses) Configure and enable switch ports as trunk ports. See the "Configuring and Enabling Switch Ports as Trunk Ports" section on page 9-19.

Step 4 Complete the interface configuration according to the "Completing Interface Configuration (All Models)" section on page 9-20.

Configuring VLAN Interfaces

This section describes how to configure VLAN interfaces. For more information about ASA 5505 interfaces, see the "ASA 5505 Interfaces" section on page 9-2.

Detailed Steps

Note 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.

Step 1 Choose the **Configuration** > **Device Setup** > **Interfaces** pane.

Step 2 On the Interfaces tab, click **Add**.

The Add Interface dialog box appears with the General tab selected.

Step 3 In the Available Switch Ports pane, choose a switch port, and click **Add**.

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.

Repeat for any other switch ports that you want to carry this VLAN.

Note Removing a switch port from an interface essentially just reassigns that switch port to VLAN 1, because the default VLAN interface for switch ports is VLAN 1.

Step 4 Click the **Advanced** tab.

Note You receive an error message about setting the IP address. You can either set the IP address and other parameters now, or you can finish configuring the VLAN and switch ports by clicking **Yes**, and later set the IP address and other parameters according to the "Completing Interface Configuration (All Models)" section on page 9-20.

Step 5 In the VLAN ID field, enter the VLAN ID for this interface, between 1 and 4090.

If you do not want to assign the VLAN ID, ASDM assigns one for you randomly.

Step 6 (Optional for the Base license) To allow this interface to be the third VLAN by limiting it from initiating contact to one other VLAN, in the Block Traffic From this Interface to drop-down list, choose the VLAN to which this VLAN interface cannot initiate traffic.

With the Base license, you can only configure a third VLAN if you use this command 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 this 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 configure this setting before setting the name on the third interface; the security appliance does not allow three fully functioning VLAN interfaces with the Base license on the ASA 5505 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.

To configure the MAC address and MTU, see the "Configuring Advanced Interface Parameters" section on page 9-25.

Step 7 Click **OK**.

What to Do Next

Configure the switch ports. See the "Configuring and Enabling Switch Ports as Access Ports" section on page 9-17 and the "Configuring and Enabling Switch Ports as Trunk Ports" section on page 9-19.

Configuring and Enabling Switch Ports as Access Ports

By default (with no configuration), all switch ports are shut down, and assigned to VLAN 1. To assign a switch port to a single VLAN, configure it as an access port. To create a trunk port to carry multiple VLANs, see the "Configuring and Enabling Switch Ports as Trunk Ports" section on page 9-19. If you have a factory default configuration, see the "ASA 5505 Default Configuration" section on page 1-2 to check if you want to change the default interface settings according to this procedure.

For more information about ASA 5505 interfaces, see the "ASA 5505 Interfaces" section on page 9-2.

Caution The ASA 5505 security appliance does not support Spanning Tree Protocol for loop detection in the network. Therefore you must ensure that any connection with the security appliance does not end up in a network loop.

Detailed Steps

Step 1 Choose the **Configuration** > **Device Setup** > **Interfaces** pane.

Step 2 Click the **Switch Ports** tab.

Step 3 Click the switch port you want to edit.

The Edit Switch Port dialog box appears.

Step 4 To enable the switch port, check the **Enable SwitchPort** check box.

Step 5 In the Mode and VLAN IDs area, click the **Access** radio button.

Step 6 In the VLAN ID field, enter the VLAN ID associated with this switch port. The VLAN ID can be between 1 and 4090.

By default, the VLAN ID is derived from the VLAN interface configuration you completed in "Configuring VLAN Interfaces" section on page 9-16 (on the Configuration > Device Setup > Interfaces > Interfaces > Add/Edit Interface dialog box). You can change the VLAN assignment in this dialog box. Be sure to apply the change to update the VLAN configuration with the new information. If you want to specify a VLAN that has not yet been added, we suggest you add the VLAN according to the "Configuring VLAN Interfaces" section on page 9-16 rather than specifying it in this dialog box; in either case, you need to add the VLAN according to the "Configuring VLAN Interfaces" section on page 9-16 and assign the switch port to it.

Step 7 (Optional) To prevent the switch port from communicating with other protected switch ports on the same VLAN, check the **Isolated** check box.

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.

Step 8 (Optional) From the Duplex drop-down list, choose **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.

Step 9 (Optional) From the Speed drop-down list, choose **10**, **100**, or **Auto**.

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.

Step 10 Click **OK**.

What to Do Next

If you want to configure a switch port as a trunk port, see the "Configuring and Enabling Switch Ports as Trunk Ports" section on page 9-19.

To complete the interface configuration, see the "Completing Interface Configuration (All Models)" section on page 9-20.

Configuring and Enabling Switch Ports as Trunk Ports

This procedure tells how to create a trunk port that can carry multiple VLANs using 802.1Q tagging. Trunk mode is available only with the Security Plus license.

To create an access port, where an interface is assigned to only one VLAN, see the "Configuring and Enabling Switch Ports as Access Ports" section on page 9-17.

For more information about ASA 5505 interfaces, see the "ASA 5505 Interfaces" section on page 9-2.

Detailed Steps

Step 1 Choose the **Configuration** > **Device Setup** > **Interfaces** pane.

Step 2 Click the **Switch Ports** tab.

Step 3 Click the switch port you want to edit.

The Edit Switch Port dialog box appears.

Step 4 To enable the switch port, check the **Enable SwitchPort** check box.

Step 5 In the Mode and VLAN IDs area, click the **Trunk** radio button.

Step 6 In the VLAN IDs field, enter the VLAN IDs associated with this switch port, separated by commas. The VLAN ID can be between 1 and 4090.

You can include the native VLAN in this field, but it is not required; the native VLAN is passed whether it is included in this field or not.

This switch port cannot pass traffic until you assign at least one VLAN to it, native or non-native.

If the VLANs are already in your configuration, after you apply the change, the Configuration > Device Setup > 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 according to the "Configuring VLAN Interfaces" section on page 9-16 rather than specifying it in this dialog box; in either case, you need to add the VLAN according to the "Configuring VLAN Interfaces" section on page 9-16 and assign the switch port to it.

Step 7 To configure the native VLAN, check the **Configure Native VLAN** check box, and enter the VLAN ID in the Native VLAN ID field. The VLAN ID can be between 1 and 4090.

Packets on the native VLAN are not modified when sent over the trunk. For example, if a port has VLANs 2, 3 and 4 assigned to it, and VLAN 2 is the native VLAN, then packets on VLAN 2 that egress the port are not modified with an 802.1Q header. Frames which ingress (enter) this port and have no 802.1Q header are put into VLAN 2.

Each port can only have one native VLAN, but every port can have either the same or a different native VLAN.

Step 8 (Optional) To prevent the switch port from communicating with other protected switch ports on the same VLAN, check the **Isolated** check box.

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.

Step 9 (Optional) From the Duplex drop-down list, choose **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.

Step 10 (Optional) From the Speed drop-down list, choose **10**, **100**, or **Auto**.

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.

Step 11 Click **OK**.

What to Do Next

To complete the interface configuration, see the "Completing Interface Configuration (All Models)" section on page 9-20.

Completing Interface Configuration (All Models)

This section includes tasks to complete the interface configuration for all models.

Note For multiple context mode, complete the tasks in this section in the context execution space. In the Configuration > Device List pane, double-click the context name under the active device IP address.

This section includes the following topics:

- Configuring General Interface Parameters, page 9-21
- Configuring Advanced Interface Parameters, page 9-25
- Configuring IPv6 Addressing, page 9-26
- Configuring the Link-Local Address on an Interface (Transparent Firewall Mode), page 9-30

Task Flow for Completing Interface Configuration

Step 1 Complete the procedures in the "Starting Interface Configuration (ASA 5510 and Higher)" section on page 9-8 or the "Starting Interface Configuration (ASA 5505)" section on page 9-15.

Step 2 (Multiple context mode) In the Configuration > Device List pane, double-click the context name under the active device IP address.

Step 3 Configure general interface parameters, including the interface name, security level, and IPv4 address. See the "Configuring General Interface Parameters" section on page 9-21.

For transparent mode, you do not configure IP addressing per interface, except for the management-only interface (see the "Information About the Management Interface" section on page 9-22). You do need to configure the other parameters in this section, however. To set the global management address for transparent mode, see the "Configuring the Management IP Address for Transparent Firewall Mode" section on page 8-1.

Step 4 (Optional) Configure the MAC address and the MTU. See the "Configuring Advanced Interface Parameters" section on page 9-25.

Step 5 (Optional) Configure IPv6 addressing. See the "Configuring IPv6 Addressing" section on page 9-26

For transparent mode, you do not configure IP addressing per interface, except for the management-only interface (see the "Information About the Management Interface" section on page 9-22). This section includes how to set the link-local address in transparent mode, but this task is usually not required. To set the global management address for transparent mode, see the "Configuring the Management IP Address for Transparent Firewall Mode" section on page 8-1.

Configuring General Interface Parameters

This procedure describes how to set the name, security level, IPv4 address and other options.

For the ASA 5510 and higher, you must configure interface parameters for the following interface types:

- Physical interfaces
- VLAN subinterfaces
- Redundant interfaces

For the ASA 5505, you must configure interface parameters for the following interface types:

- VLAN interfaces

Guidelines and Limitations

- For the ASA 5550 security appliance, for maximum throughput, be sure to balance your traffic over the two interface slots; for example, assign the inside interface to slot 1 and the outside interface to slot 0.
- For information about security levels, see the "Security Levels" section on page 9-5.
- If you are using failover, do not use this procedure to name interfaces that you are reserving for failover and Stateful Failover communications. See the "Failover Link Configuration" section on page 17-10 and "State Link Configuration" section on page 17-11 to configure the failover and state links.
- In routed firewall mode, set the IP address for all interfaces.
- In transparent firewall mode, do not set the IP address for each interface, but rather set it for the whole security appliance or context. The exception is for the Management 0/0 or 0/1 management-only interface, which does not pass through traffic. To set the transparent firewall mode whole security appliance or context management IP address, see the "Configuring the Management IP Address for Transparent Firewall Mode" section on page 8-1. To set the IP address of the Management 0/0 or 0/1 interface or subinterface, use this procedure.

Restrictions

PPPoE is not supported in multiple context mode or transparent firewall mode.

Information About the Management Interface

The ASA 5510 and higher security appliance includes a dedicated management interface called Management 0/0 or Management 0/1, depending on your model, which is meant to support traffic to the security appliance. However, you can configure any interface to be a management-only interface. Also, for Management 0/0 or 0/1, you can disable management-only mode so the interface can pass through traffic just like any other interface.

Transparent firewall mode allows only two interfaces to pass through traffic; however, on the ASA 5510 and higher security appliance, you can use the Management 0/0 or 0/1 interface (either the physical interface or a subinterface) as a third interface for management traffic. The mode is not configurable in this case and must always be management-only.

Prerequisites

- Complete the procedures in the “Starting Interface Configuration (ASA 5510 and Higher)” section on page 9-8 or the “Starting Interface Configuration (ASA 5505)” section on page 9-15.
- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Detailed Steps

Step 1 Choose the **Configuration** > **Device Setup** > **Interfaces** pane.

For the ASA 5505, the Interfaces tab shows by default.

Step 2 Choose the interface row, and click **Edit**.

The Edit Interface dialog box appears with the General tab selected.

Step 3 In the Interface Name field, enter a name up to 48 characters in length.

Step 4 In the Security level field, enter a level between 0 (lowest) and 100 (highest).

See the “Security Levels” section on page 9-5 for more information.

Step 5 (Optional) To set this interface as a management-only interface, check the **Dedicate this interface to management-only** check box.

Through traffic is not accepted on a management-only interface. For the ASA 5510 and higher, see the “Information About the Management Interface” section on page 9-22 for more information.

Step 6 If the interface is not already enabled, check the **Enable Interface** check box.

Step 7 To set the IP address, one of the following options.

Note For use with failover, you must set the IP address and standby address manually; DHCP and PPPoE are not supported. Set the standby IP addresses on the Configuration > Device Management > High Availability > Failover > Interfaces tab

In transparent firewall mode, do not set the IP address for each interface, but rather set it for the whole security appliance or context. The exception is for the Management 0/0 or 0/1 management-only interface, which does not pass through traffic.

- To set the IP address manually, click the **Use Static IP** radio button and enter the IP address and mask.
- To obtain an IP address from a DHCP server, click the **Obtain Address via DHCP** radio button.
  a. (Optional) To force a MAC address to be stored inside a DHCP request packet for option 61 instead of the default internally-generated string, click the **Use MAC Address** radio button.

     Some ISPs expect option 61 to be the interface MAC address. If the MAC address is not included in the DHCP request packet, then an IP address will not be assigned.

  b. (Optional) To obtain the default route from the DHCP server, check **Obtain Default Route Using DHCP**.
  c. (Optional) To assigns an administrative distance to the learned route, enter a value between 1 and 255 in the DHCP Learned Route Metric field. If this field is left blank, the administrative distance for the learned routes is 1.
  d. (Optional) To enable tracking for DHCP-learned routes, check **Enable Tracking for DHCP Learned Routes**. Set the following values:

     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.

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

     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.

  e. (Optional) To renew the lease, click **Renew DHCP Lease**.
- (Single, routed mode only) To obtain an IP address using PPPoE, check **Use PPPoE**.
  a. In the Group Name field, specify a group name.
  b. In the PPPoE Username field, specify the username provided by your ISP.
  c. In the PPPoE Password field, specify the password provided by your ISP.
  d. In the Confirm Password field, retype the password.
  e. For PPP authentication, click either the **PAP**, **CHAP**, or **MSCHAP** radio button.

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.

f. (Optional) To store the username and password in Flash memory, check the **Store Username and Password in Local Flash** check box.

The security appliance stores the username and password in a special location of NVRAM. If an Auto Update Server sends a **clear configure** 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.

g. (Optional) To display the PPPoE IP Address and Route Settings dialog box where you can choose addressing and tracking options, click **IP Address and Route Settings**. See the "PPPoE IP Address and Route Settings" section on page 9-24 for more information.

Step 8 (Optional) In the Description field, enter a description for this interface.

The description can be up to 240 characters on a single line, without carriage returns. 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.

Note (ASA 5510 and higher) For information about the Configure Hardware Properties button, see the "Enabling the Physical Interface and Configuring Ethernet Parameters" section on page 9-9.

Step 9 Click **OK**.

What to Do Next

- (Optional) Configure the MAC address and MTU. See the "Configuring Advanced Interface Parameters" section on page 9-25.
- (Optional) Configure IPv6 addressing. See the "Configuring IPv6 Addressing" section on page 9-26

PPPoE IP Address and Route Settings

The Configuration > Interfaces > Add/Edit Interface > General > PPPoE IP Address and Route Settings > PPPoE IP Address and Route Settings dialog box 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 administrative distance for the learned routes is 1.
  - Enable tracking—Check this check box 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.

Configuring Advanced Interface Parameters

This section describes how to configure MAC addresses for interfaces and how to set the MTU.

Information About 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. A redundant interface uses the MAC address of the first physical interface that you add. If you change the order of the member interfaces in the configuration, then the MAC address changes to match the MAC address of the interface that is now listed first. If you assign a MAC address to the redundant interface using this command, then it is used regardless of the member interface MAC addresses.

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 11-2 for more information. You can assign each MAC address manually, or you can automatically generate MAC addresses for shared interfaces in contexts. See the "Automatically Assigning MAC Addresses" section on page 11-18 to automatically generate MAC addresses. If you automatically generate MAC addresses, you can use this procedure 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.

Detailed Steps

Step 1 Choose the **Configuration** > **Device Setup** > **Interfaces** pane.

For the ASA 5505, the Interfaces tab shows by default.

Step 2 Choose the interface row, and click **Edit**.

The Edit Interface dialog box appears with the General tab selected.

Step 3 Click the **Advanced** tab.

Step 4 To set the MTU or to enable jumbo frame support (ASA 5580 only), enter the value in the MTU field, between 300 and 65,535 bytes.

The default is 1500 bytes.

- For the ASA 5580 in single mode—If you enter a value for any interface that is greater than 1500, then you enable jumbo frame support automatically for all interfaces. If you set the MTU for all interfaces back to a value under 1500, then jumbo frame support is disabled.
- For the ASA 5580 in multiple mode—If you enter a value for any interface that is greater than 1500, then be sure to enable jumbo frame support in the system configuration. See the "Enabling Jumbo Frame Support (ASA 5580, Multiple Mode)" section on page 9-32.

Note Enabling or disabling jumbo frame support requires you to reboot the security appliance.

A jumbo frame is an Ethernet packet larger than the standard maximum of 1518 bytes (including Layer 2 header and FCS), up to 9216 bytes. Jumbo frames require extra memory to process, and assigning more memory for jumbo frames might limit the maximum use of other features, such as access lists.

Step 5 To manually assign a MAC address to this interface, enter a MAC address in the Active Mac Address field 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.

Step 6 If you use failover, enter the standby MAC address in the Standby Mac Address field. 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.

What to Do Next

(Optional) Configure IPv6 addressing. See the "Configuring IPv6 Addressing" section on page 9-26

Configuring IPv6 Addressing

This section describes how to configure IPv6 addressing. For more information about IPv6, see the "IPv6 Addresses" section on page B-5.

For transparent mode, use this section for the Management 0/0 or 0/1 interface. To configure the global IPv6 management address for transparent mode, see the "Configuring the Management IP Address for Transparent Firewall Mode" section on page 8-1. If you do not configure a management address, you can configure the link-local addresses in transparent mode according to the "Configuring the Link-Local Address on an Interface (Transparent Firewall Mode)" section on page 9-30.

Information About IPv6 Addressing

When you configure an IPv6 address on an interface, you can assign one or several IPv6 addresses to the interface at one time, such as an IPv6 link-local address and a global address. However, at a minimum, you must configure a link-local address.

Every IPv6-enabled interface must include at least one link-local address. When you configure a global address, a link-local addresses is automatically configured on the interface, so you do not also need to specifically configure a link-local address. These link-local addresses can only be used to communicate with other hosts on the same physical link.

When IPv6 is used over Ethernet networks, the Ethernet MAC address can be used to generate the 64-bit interface ID for the host. This is called the EUI-64 address. Because MAC addresses use 48 bits, additional bits must be inserted to fill the 64 bits required. The last 64 bits are used for the interface ID. For example, FE80::/10 is a link-local unicast IPv6 address type in hexadecimal format.

Information About Duplicate Address Detection

During the stateless autoconfiguration process, duplicate address detection (DAD) verifies the uniqueness of new unicast IPv6 addresses before the addresses are assigned to interfaces (the new addresses remain in a tentative state while duplicate address detection is performed). Duplicate address detection is performed first on the new link-local address. When the link local address is verified as unique, then duplicate address detection is performed all the other IPv6 unicast addresses on the interface.

Duplicate address detection is suspended on interfaces that are administratively down. While an interface is administratively down, the unicast IPv6 addresses assigned to the interface are set to a pending state. An interface returning to an administratively up state restarts duplicate address detection for all of the unicast IPv6 addresses on the interface.

When a duplicate address is identified, the state of the address is set to DUPLICATE, the address is not used, and the following error message is generated:

```
%PIX|ASA-4-325002: Duplicate address ipv6_address/MAC_address on interface
```

If the duplicate address is the link-local address of the interface, the processing of IPv6 packets is disabled on the interface. If the duplicate address is a global address, the address is not used. However, all configuration commands associated with the duplicate address remain as configured while the state of the address is set to DUPLICATE.

If the link-local address for an interface changes, duplicate address detection is performed on the new link-local address and all of the other IPv6 address associated with the interface are regenerated (duplicate address detection is performed only on the new link-local address).

The security appliance uses neighbor solicitation messages to perform duplicate address detection. By default, the number of times an interface performs duplicate address detection is 1.

Information About Modified EUI-64 Interface IDs

RFC 3513: Internet Protocol Version 6 (IPv6) Addressing Architecture requires that the interface identifier portion of all unicast IPv6 addresses, except those that start with binary value 000, be 64 bits long and be constructed in Modified EUI-64 format. The security appliance can enforce this requirement for hosts attached to the local link.

When this command is enabled on an interface, the source addresses of IPv6 packets received on that interface are verified against the source MAC addresses to ensure that the interface identifiers use the Modified EUI-64 format. If the IPv6 packets do not use the Modified EUI-64 format for the interface identifier, the packets are dropped and the following system log message is generated:

```
%PIX|ASA-3-325003: EUI-64 source address check failed.
```

The address format verification is only performed when a flow is created. Packets from an existing flow are not checked. Additionally, the address verification can only be performed for hosts on the local link. Packets received from hosts behind a router will fail the address format verification, and be dropped, because their source MAC address will be the router MAC address and not the host MAC address.

Restrictions

The security appliance does not support IPv6 anycast addresses.

Detailed Steps

Step 1 Choose the **Configuration** > **Device Setup** > **Interfaces** pane.

Step 2 Choose an interface, and click **Edit**.

The Edit Interface dialog box appears with the General tab selected.

Step 3 Click the **IPv6** tab.

Step 4 (Optional) To enforce the use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link, check the **Enforce EUI-64** check box.

If the interface identifiers do not conform to the modified EUI-64 format, an error message appears. See the "Information About Modified EUI-64 Interface IDs" section on page 9-28 for more information.

Step 5 Configure the global IPv6 address using one of the following methods.

Note If you do not want to configure a global IPv6 address, you can configure the link-local addresses either automatically by checking the **Enable IPv6** check box, or manually by entering a value in the Link-local address field in the Interface IPv6 Addresses area. A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. See the "IPv6 Addresses" section on page B-5 for more information about IPv6 addressing.

If you configure a global IPv6 address (or manually configure a link-local address), checking or unchecking the **Enable IPv6** check box does not affect how IPv6 operates; IPv6 continues to be enabled.

- Stateless autoconfiguration—In the Interface IPv6 Addresses area, check the **Enable address autoconfiguration** check box.

Enabling stateless autconfiguration on the interface configures IPv6 addresses based upon prefixes received in Router Advertisement messages. A link-local address, based on the Modified EUI-64 interface ID, is automatically generated for the interface when stateless autoconfiguration is enabled.

- Manual configuration—To manually configure a global IPv6 address:
  a. In the Interface IPv6 Addresses area, click **Add**.

     The Add IPv6 Address for Interface dialog box appears.

  b. In the Address/Prefix Length field, enter the global IPv6 address and the IPv6 prefix length. For example, 2001:0DB8::BA98:0:3210/48. See the "IPv6 Addresses" section on page B-5 for more information about IPv6 addressing.
  c. (Optional) To use the Modified EUI-64 interface ID in the low order 64 bits of the address, check the **EUI-64** check box.
  d. Click **OK**.

Step 6 (Optional) In the top area, customize the IPv6 configuration by configuring the following options:

- DAD Attempts—This setting configures the number of consecutive neighbor solicitation messages that are sent on an interface while DAD is performed on IPv6 addresses. Valid values are from 0 to 600. A zero value disables DAD processing on the specified interface. The default is one message.
- NS Interval—Enter the neighbor solicitation message interval. The neighbor solicitation message requests the link-layer address of a target node. Valid values are from 1000 to 3600000 milliseconds. The default is 1000 milliseconds.
- Reachable Time—Enter the amount of time in seconds that a remote IPv6 node is considered reachable after a reachability confirmation event has occurred. Valid values are from 0 to 3600000 milliseconds. The default is zero. A configured time enables the detection of unavailable neighbors. Shorter times enable detection more quickly; however, very short configured times are not recommended in normal IPv6 operation.
- RA Lifetime—Enter the amount of time that IPv6 router advertisement transmissions are considered valid. Valid values are from 0 to 9000 seconds. The default is 1800 seconds. Router advertisement transmissions include a preference level and a lifetime field for each advertised router address. These transmissions provide route information and indicate that the router is still operational to network hosts.
- RA Interval—Enter the interval between IPv6 router advertisement transmissions. Valid values are from 3 to 1800 seconds. The default is 200 seconds. To list the router advertisement transmission interval in milliseconds, check the **RA Interval in Milliseconds** check box. Valid values are from 500 to 1800000 milliseconds.
- To allow the generation of addresses for hosts, make sure that the Suppress RA check box is unchecked. This is the default setting if IPv6 unicast routing is enabled. To prevent the generation of IPv6 router advertisement transmissions, check the **Suppress RA** check box.

Step 7 (Optional) To configure which IPv6 prefixes are included in IPv6 router advertisements, complete the following.

By default, prefixes configured as addresses on an interface are advertised in router advertisements. If you configure prefixes for advertisement using this area, then only these prefixes are advertised.

a. In the Interface IPv6 Prefixes area, click **Add**.

   The Add IPv6 Prefix for Interface dialog box appears.

b. In the Address/Prefix Length field, enter the IPv6 address with the prefix length. To configure settings that apply to all prefixes, check the **Default Values** check box instead of entering an Address.

c. (Optional) To indicate that the IPv6 prefix is not advertised, check the **No Advertisements** check box.

d. (Optional) To indicate that the specified prefix is not used for on-link determination, check the **Off-link** check box.

e. (Optional) To indicate to hosts on the local link that the specified prefix cannot be used for IPv6 autoconfiguration, check the **No Auto-Configuration** check box.

f. In the Prefix Lifetime area, choose one of the following:

   – Lifetime Duration—Specify the following:

   A valid lifetime for the prefix in seconds from the drop-down list. This setting is the amount of time that the specified IPv6 prefix is advertised as being valid. The maximum value represents infinity. Valid values are from 0 to 4294967295. The default is 2592000 (30 days).

   A preferred lifetime for the prefix from the drop-down list. This setting is the amount of time that the specified IPv6 prefix is advertised as being preferred. The maximum value represents infinity. Valid values are from 0 to 4294967295. The default setting is 604800 (seven days).

   – Lifetime Expiration Date—Specify the following:

   Choose a valid month and day from the drop-down list, and then enter a time in hh:mm format.

   Choose a preferred month and day from the drop-down list, and then enter a time in hh:mm format.

Step 8 Click **OK**.

You return to the Edit Interface dialog box.

Step 9 Click **OK**.

You return to the Configuration > Device Setup > Interfaces pane.

Configuring the Link-Local Address on an Interface (Transparent Firewall Mode)

If you only need to configure a link-local address and are not going to assign any other IPv6 addresses, you have the option of manually defining the link-local address.

To assign a link-local address to an interface, perform the following steps:

Step 1 Choose the **Configuration > Device Setup > Interfaces** pane.

Step 2 Select an interface, and click **Edit**.

The Edit Interface dialog box appears with the General tab selected.

Step 3 Click the **IPv6** tab.

Step 4 (Optional) To enforce the use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link, check the **Enforce EUI-64** check box.

If the interface identifiers do not conform to the modified EUI-64 format, an error message appears. See the "Information About Modified EUI-64 Interface IDs" section on page 9-28 for more information.

Step 5 To set the link-local address, enter an address in the Link-local address field.

A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. See the "IPv6 Addresses" section on page B-5 for more information about IPv6 addressing.

Step 6 Click **OK**.

Allowing Same Security Level Communication

By default, interfaces on the same security level cannot communicate with each other, and packets cannot enter and exit the same interface. This section describes how to enable inter-interface communication when interfaces are on the same security level, and how to enable intra-interface communication.

Information About Inter-Interface Communication

Allowing interfaces on the same security level to communicate with each other 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.

If you enable same security interface communication, you can still configure interfaces at different security levels as usual.

Note If you enable NAT control, you do not need to configure NAT between same security level interfaces. See the "NAT and Same Security Level Interfaces" section on page 22-12 for more information on NAT and same security level interfaces.

Information About Intra-Interface Communication

Intra-interface communication might be useful for VPN traffic that enters an interface, but is then routed out the same interface. The VPN traffic might be unencrypted in this case, or it might be reencrypted for another VPN connection. For example, if you have a hub and spoke VPN network, where the security appliance is the hub, and remote VPN networks are spokes, for one spoke to communicate with another spoke, traffic must go into the security appliance and then out again to the other spoke.

Note All traffic allowed by this feature is still subject to firewall rules. Be careful not to create an asymmetric routing situation that can cause return traffic not to traverse the security appliance.

Restrictions

This feature is only available in routed firewall mode.

Detailed Steps

To disable these settings, use the **no** form of the command.

- To enable interfaces on the same security level to communicate with each other, from the Configuration > Interfaces pane, check **Enable traffic between two or more interfaces which are configured with same security level**.
- To enable communication between hosts connected to the same interface, check **Enable traffic between two or more hosts connected to the same interface**.

Enabling Jumbo Frame Support (ASA 5580, Multiple Mode)

A jumbo frame is an Ethernet packet larger than the standard maximum of 1518 bytes (including Layer 2 header and FCS), up to 9216 bytes. You can enable support for jumbo frames for all interfaces by increasing the amount of memory to process Ethernet frames. Assigning more memory for jumbo frames might limit the maximum use of other features, such as access lists.

Note Other platform models do not support jumbo frames.

Prerequisites

In multiple context mode, set this option in the system execution space. In single mode, setting the MTU larger than 1500 bytes automatically enables jumbo frames.

Detailed Steps

To enable jumbo frame support, choose the **Configuration > Context Management > Interfaces pane**, and click the **Enable jumbo frame support** check box.

Note Changes in this setting require you to reboot the security appliance.

Be sure to set the MTU for each interface that needs to transmit jumbo frames to a higher value than the default 1500; for example, set the value to 9000. See the "Configuring Advanced Interface Parameters" section on page 9-25. Set the MTU within each context.

Monitoring Interfaces

To monitor interfaces, see Chapter 42, "Monitoring Interfaces."

Feature History for Interfaces

Table 9-1 lists the release history for this feature.

Table 9-1 Feature History for Interfaces

Feature Name	Releases	Feature Information
Increased VLANs	7.0(5)	Increased the following limits: • ASA5510 Base license VLANs from 0 to 10. • ASA5510 Security Plus license VLANs from 10 to 25. • ASA5520 VLANs from 25 to 100. • ASA5540 VLANs from 100 to 200.
Increased VLANs	7.2(2)	The maximum number of VLANs for the Security Plus license on the ASA 5505 security appliance was increased from 5 (3 fully functional; 1 failover; one restricted to a backup interface) to 20 fully functional interfaces. In addition, the number of trunk ports was increased from 1 to 8. Now there are 20 fully functional interfaces, you do not need to use the backup interface command to cripple a backup ISP interface; you can use a fully-functional interface for it. The backup interface command is still useful for an Easy VPN configuration. VLAN limits were also increased for the ASA 5510 security appliance (from 10 to 50 for the Base license, and from 25 to 100 for the Security Plus license), the ASA 5520 security appliance (from 100 to 150), the ASA 5550 security appliance (from 200 to 250).
Gigabit Ethernet Support for the ASA 5510 Security Plus License	7.2(3)	The ASA 5510 security appliance now supports GE (Gigabit Ethernet) for port 0 and 1 with the Security Plus license. If you upgrade the license from Base to Security Plus, the capacity of the external Ethernet0/0 and Ethernet0/1 ports increases from the original FE (Fast Ethernet) (100 Mbps) to GE (1000 Mbps). The interface names will remain Ethernet 0/0 and Ethernet 0/1. Use the **speed** command to change the speed on the interface and use the **show interface** command to see what speed is currently configured for each interface.
Native VLAN support for the ASA 5505	7.2(4)/8.0(4)	You can now include the native VLAN in an ASA 5505 trunk port using the **switchport trunk native vlan** command.

Table 9-1 Feature History for Interfaces (continued)

Feature Name	Releases	Feature Information
Gigabit Ethernet Support for the ASA 5510 Base License	7.2(4)/8.0(4)	The ASA 5510 security appliance now supports GE (Gigabit Ethernet) for port 0 and 1 in the Base license (support was previously added for the Security Plus license). The capacity of the external Ethernet0/0 and Ethernet0/1 ports increases from the original FE (Fast Ethernet) (100 Mbps) to GE (1000 Mbps). The interface names will remain Ethernet 0/0 and Ethernet 0/1. Use the **speed** command to change the speed on the interface and use the **show interface** command to see what speed is currently configured for each interface.
Jumbo packet support for the ASA 5580	8.1(1)	The Cisco ASA 5580 supports jumbo frames when you enter the **jumbo-frame reservation** command. A jumbo frame is an Ethernet packet larger than the standard maximum of 1518 bytes (including Layer 2 header and FCS), up to 9216 bytes. You can enable support for jumbo frames for all interfaces by increasing the amount of memory to process Ethernet frames. Assigning more memory for jumbo frames might limit the maximum use of other features, such as access lists. In ASDM, see Configuration > Device Setup > Interfaces > Add/Edit Interface > Advanced.
Increased VLANs for the ASA 5580	8.1(2)	The number of VLANs supported on the ASA 5580 are increased from 100 to 250.
Support for Pause Frames for Flow Control on the ASA 5580 10 Gigabit Ethernet Interfaces	8.2(2)	You can now enable pause (XOFF) frames for flow control. The following screens were modified: (Single Mode) Configuration > Device Setup > Interfaces > Add/Edit Interface > General, (Multiple Mode, System) Configuration > Interfaces > Add/Edit Interface.

CHAPTER 10

Configuring Management Access to the SSC

This chapter describes how to configure management access to a Security Services Card (SSC) that is installed in the security appliance. The SSC runs the IPS application. For information about configuring IPS, see Configuring the IPS Application on the AIP SSM and SSC, page 30-1.

This chapter includes the following sections:

- Information About Management Access to the SSC, page 10-1
- Guidelines and Limitations, page 10-2
- Default Settings, page 10-2
- Configuring the SSC Management Interface, page 10-3
- Where to Go Next, page 10-5

Information About Management Access to the SSC

You can manage the module application using ASDM or by using the module application CLI. For information about using the CLI, see the *Cisco ASA 5500 Series Configuration Guide using the CLI*.

This section includes the following topics:

- Information About Using ASDM to Manage the SSC, page 10-1
- Other Uses for the SSC Management Interface, page 10-2
- Routing Considerations for Accessing the Management Interface, page 10-2

Information About Using ASDM to Manage the SSC

After you launch ASDM on the security appliance, ASDM connects to the SSC management interface to configure the module application. You can configure a VLAN as a management VLAN to allow access to an internal management IP address over the backplane. To change the network parameters, see the "Configuring the SSC Management Interface, page 10-3.

See the Default Settings, page 10-2 for information about the default management interface parameters.

Other Uses for the SSC Management Interface

The module management interface can be used for sending syslog messages or allowing updates for the module application, such as signature database updates on the AIP SSC.

Routing Considerations for Accessing the Management Interface

To make sure ASDM can manage the SSC, be sure that the security appliance can access the module management interface address.

Be sure to configure an IP address for the security appliance VLAN that you are also using for the SSC management interface, and assign that VLAN to a switch port so that the SSC interface is physically connected to the network. The SSC management interface will then be on a directly-connected network for the security appliance, so ASDM can access the management interface without any additional routing configuration.

Note If the default gateway is set to the security appliance, traffic from the SSC to devices directly connected to the security appliance go through; however, if the destination is a hop away (that is, on a different gateway), the traffic does not go through.

If you have multiple networks on an internal link, make sure that you change the default gateway to an internal router instead of leaving it at the default setting of the security appliance.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

See the chapter for each SSM or SSC application for context mode guidelines.

Firewall Mode Guidelines

See the chapter for each SSM or SSC application for firewall mode guidelines.

Failover Guidelines

Make sure you configure the management IP addresses on both units to be on the same subnet and VLAN.

Model Guidelines

For model support for the SSC, see the“SSM and SSC Support Per Model” section on page 3-2.

Additional Guidelines

You cannot set up the SSC in ASDM if you use an IP address that goes through NAT.

Default Settings

Table 10-1 lists the default network settings for SSCs.

Table 10-1 Default Network Parameters

Parameters	Default
Management VLAN	VLAN 1
Management IP address	192.168.1.2/24
Gateway	192.168.1.1

Note The default management IP address on the security appliance is 192.168.1.5/24.

Configuring the SSC Management Interface

An SSC does not have any external interfaces. You can configure a VLAN as a management VLAN to allow access to an internal management IP address over the backplane. By default, VLAN 1 is enabled for the SSC management address. You can only assign one VLAN as the SSC management VLAN. This section describes how to change the management VLAN. It also describes how to change the default management IP address, allowed hosts, and gateway. See the "Default Settings" section on page 10-2 for more information about defaults.

Prerequisites

For the VLAN you want to use for the SSC management interface, configure the switch port and VLAN interface on the ASA 5505. This is required so the SSC interface is physically connected to the network. You must create and assign VLANs to an interface before you can go to the SSC Setup pane and select one for the SSC.

Restrictions

Do not configure NAT for the management address. For initial setup with ASDM, you need to access the real address. After initial setup (where you set the the password in the SSC), you can configure NAT and supply ASDM with the translated address when you want to access the SSC on the Configuration > IPS pane.

Detailed Steps

Step 1 If you are configuring the SSC for the first time, in the ASDM main window, choose **Configuration > Device Setup > SSC Setup**.

Note If you click the **IPS** tab before you have configured the SSC, the Stop dialog box appears. Click **OK** to have ASDM redirect you to the SSC Setup pane. You must define the settings in the SSC Setup pane before you can access any part of the GUI.

Step 2 In the Management Interface area, do the following:

a. Choose the Interface VLAN from the drop-down list.

This setting allows you to manage the SSC using this VLAN.

Note The following settings are written to the SSC application configuration, not the security appliance configuration.

b. Enter the IP address.

c. Choose the subnet mask from the drop-down list.

d. Enter the default gateway IP address.

If the management station is on a directly-connected security appliance network, then set the gateway to be the ASA 5505 VLAN interface address. If the management station is on a remote network, then set the gateway to the address of an upstream router on the management VLAN.

Step 3 In the Management Access List area, do the following.

Note The following settings are written to the SSC application configuration, not the security appliance configuration.

a. Enter the IP address for the host network.

b. Choose the subnet mask from the drop-down list.

c. Click **Add** to add these settings to the Allowed Hosts/Networks list.

Note After you click **Add**, make sure you save the management settings you have just defined by clicking **Apply**. If you decide to remove these settings, continue to the next substep. Otherwise, go to Step 4.

d. To delete these settings, in the ASDM main window, click the **IPS** tab. Choose **Configuration > IPS > Sensor Setup > Allowed Hosts/Networks**. Choose the host or network that you want to remove from the list, and click **Delete**. To add new management settings, you can either click **Add** in the existing pane or return to the SSC Setup pane by choosing **Configuration > Device Setup > SSC Setup**.

Step 4 In the IPS Password area, do the following:

Note The following settings are written to the SSC application configuration, not the security appliance configuration.

a. Enter the password. The default password is “cisco.”

b. Enter the new password, and confirm the change.

Step 5 Click **Apply** to save the settings to the running configuration.

The SSC Setup completed dialog box appears only after the initial configuration.

Step 6 To complete the SSC application configuration and have ASDM go directly to the Configuration > IPS > Sensor Setup > Startup Wizard screen, do one of the following:

- Click the **IPS** button in the navigation pane.
- Click the **Configure the IPS SSC module** link.

Note If you want to change the SSC configuration settings at a later date, click the **IPS** tab.

Troubleshooting

To reset the password, choose **Tools > IPS Password Reset**. If you change the password, the Status dialog box appears, indicating that the new sensor password is being saved to the SSC application configuration.

After you have logged in and defined a new password, you do not need to log in to the IPS SSC application again. If you cannot connect to the IPS SSC application with the new password, restart ASDM and try to log in again.

If you have defined a new password and still have a existing password that is different than the new password, clear the password cache by choosing **File > Clear ASDM Password Cache**, then restart ASDM and try to log in again.

If you upgrade the ASA 5505 and add an SSC, make sure that the factory default IP address of the ASA 5505 starts at 192.168.1.5, because the factory default IP address of the SSC is 192.168.1.2. If a conflict occurs during configuration, a warning message appears.

At startup, make sure that the TFTP URL download location is valid, because if the ASA 5505 cannot detect a valid image to download, the SSC application does not start and the following error message appears:

```
Autoboot Error\System Halt
```

Note If you restart the security appliance, the SSC is not automatically restarted. For more information, see the “Managing SSMs and SSCs” section in the *Cisco ASA 5500 Series Configuration Guide using the CLI*.

Where to Go Next

See Chapter 30, “Configuring the IPS Application on the AIP SSM and SSC.”

CHAPTER 11

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, page 11-1
- Enabling or Disabling Multiple Context Mode, page 11-9
- Configuring Resource Classes, page 11-10
- Configuring Security Contexts, page 11-16

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, page 11-2
- Unsupported Features, page 11-2
- Context Configuration Files, page 11-2
- How the Security Appliance Classifies Packets, page 11-2
- Management Access to Security Contexts, page 11-8

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 routing. Multicast bridging is supported.
- Threat Detection

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 11-3

- Invalid Classifier Criteria, page 11-4
- Classification Examples, page 11-4

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 11-3
- Unique MAC Addresses, page 11-3
- NAT Configuration, page 11-3

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 Chapter 9, “Configuring Interfaces”), or you can automatically generate MAC addresses (see the “Automatically Assigning MAC Addresses” section on page 11-18).

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 11-1 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 11-1 Packet Classification with a Shared Interface using MAC Addresses

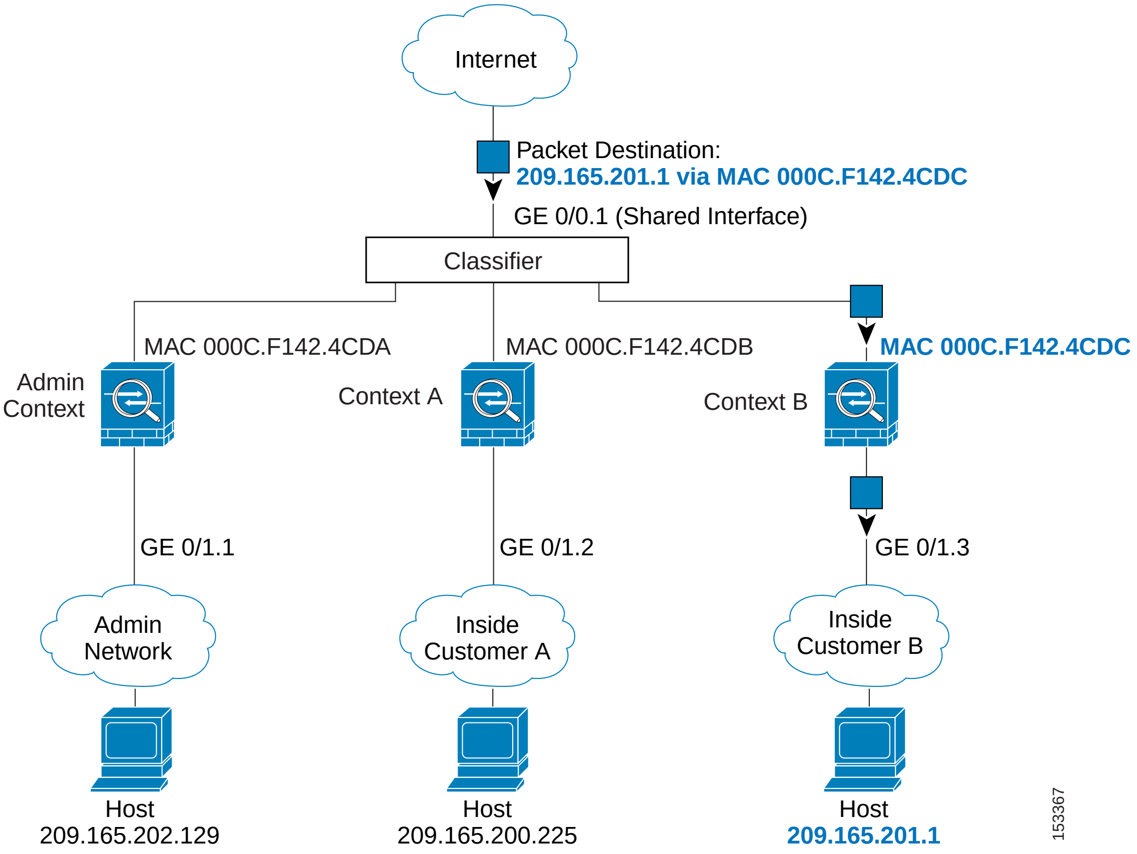

Figure 11-2 shows multiple contexts sharing an outside interface without MAC addresses assigned. The classifier assigns the packet to Context B because Context B includes the address translation that matches the destination address.

Figure 11-2 Packet Classification with a Shared Interface using NAT

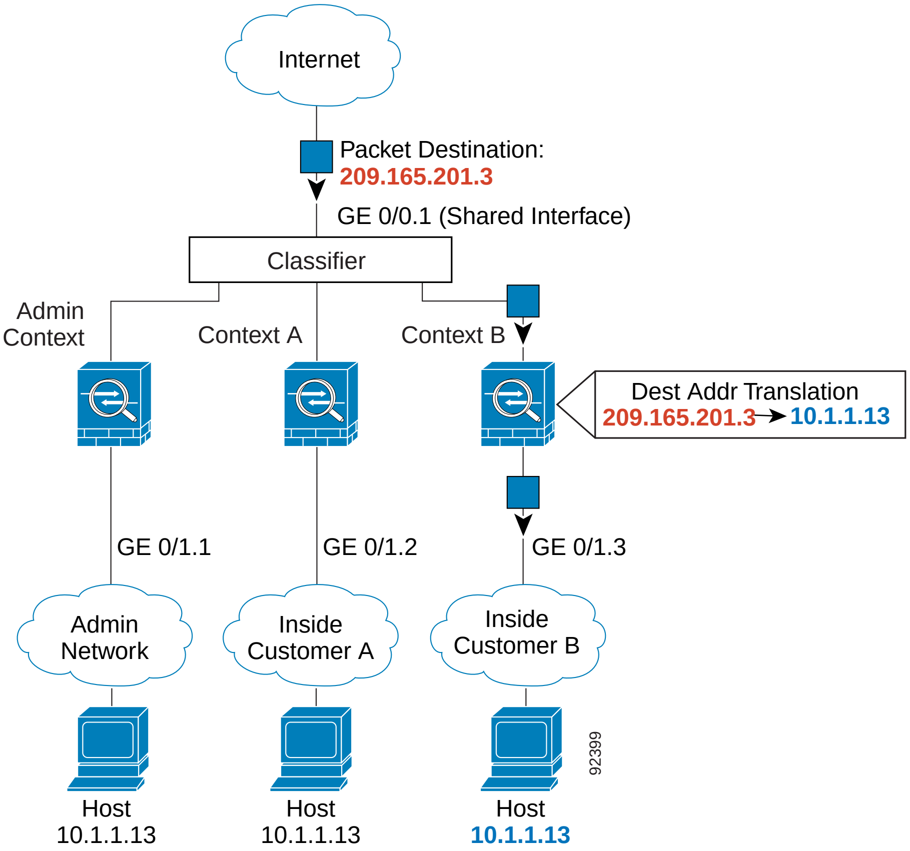

Note that all new incoming traffic must be classified, even from inside networks. Figure 11-3 shows a host on the Context B inside network accessing the Internet. The classifier assigns the packet to Context B because the ingress interface is Gigabit Ethernet 0/1.3, which is assigned to Context B.

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 11-3 Incoming Traffic from Inside Networks

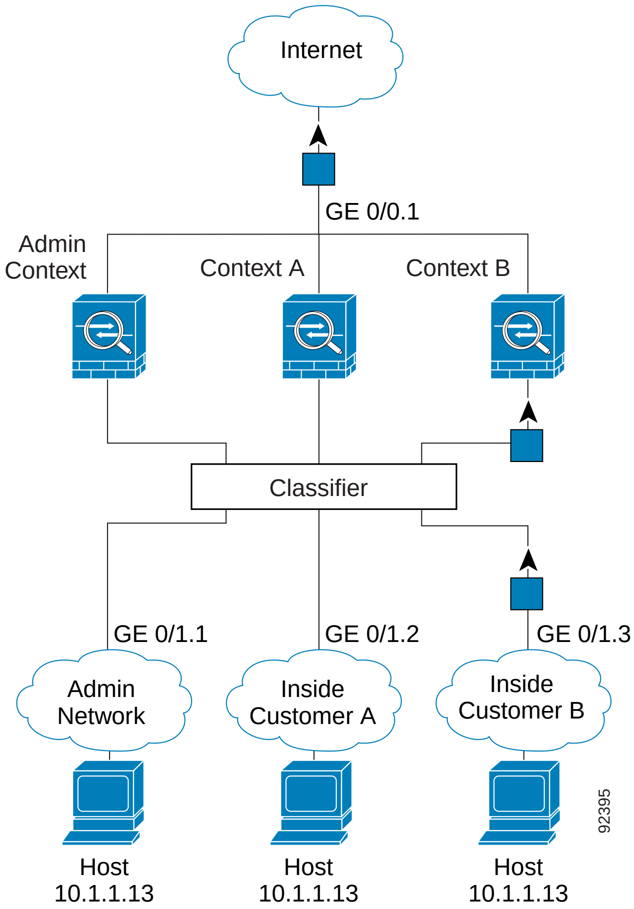

For transparent firewalls, you must use unique interfaces. Figure 11-4 shows a host on the Context B inside network accessing the Internet. The classifier assigns the packet to Context B because the ingress interface is Gigabit Ethernet 1/0.3, which is assigned to Context B.

Figure 11-4 Transparent Firewall Contexts

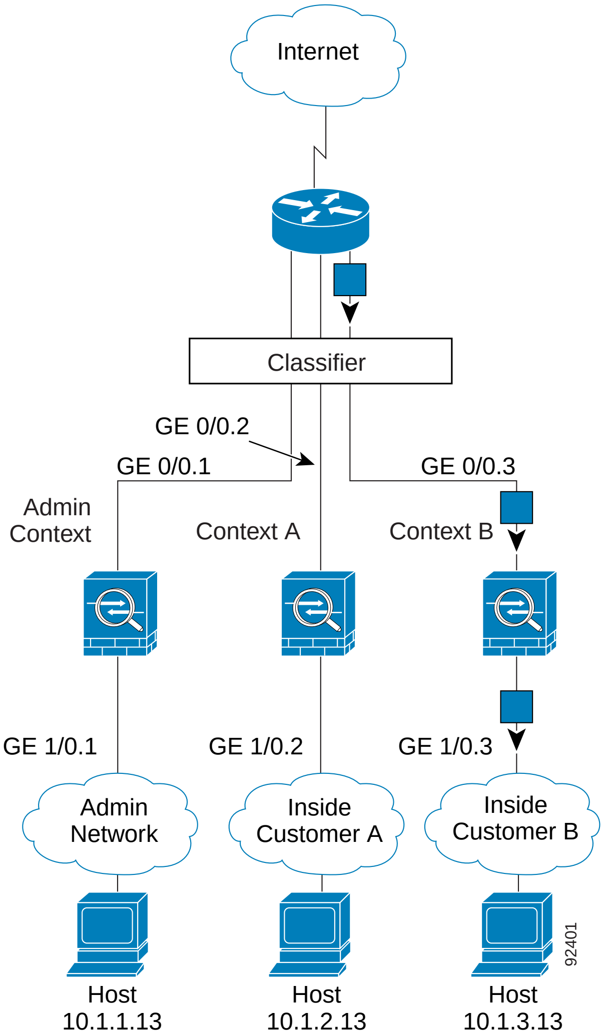

Cascading Security Contexts

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 11-5 shows a gateway context with two contexts behind the gateway.

Figure 11-5 Cascading Contexts

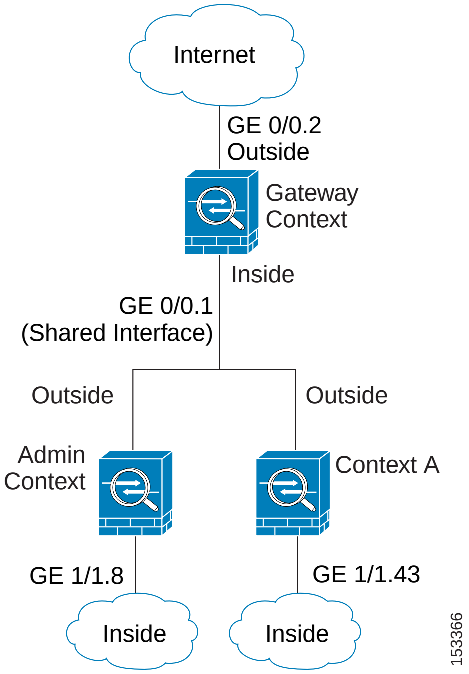

Management Access to Security Contexts

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 11-8
- Context Administrator Access, page 11-9

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 Configuring Authentication for Network Access, page 24-1 to enable Telnet, SSH, and ASDM 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 Configuring Authentication for Network Access, page 24-1 to enable Telnet, SSH, and SDM access and to configure management authentication.

Enabling or Disabling Multiple Context Mode

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 supports changing modes from single to multiple mode if you use the High Availability and Scalability Wizard and you enable Active/Active failover. See the "Accessing and Using the High Availability and Scalability Wizard" section on page 17-5 for more information.

If you do not want to use Active/Active failover or want to change back to single mode, you must change modes at the CLI. This section describes changing modes at the CLI, and includes the following topics:

- Backing Up the Single Mode Configuration, page 11-9
- Enabling Multiple Context Mode, page 11-9
- Restoring Single Context Mode, page 11-10

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
```

Note Be sure that you do not save the current running configuration, or it will overwrite the one you just copied.

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, page 11-11
- Adding a Resource Class, page 11-13
- Monitoring Context Resource Usage, page 11-15

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, page 11-11
- Default Class, page 11-12
- Class Members, page 11-13

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 11-6.)

Figure 11-6 Resource Oversubscription

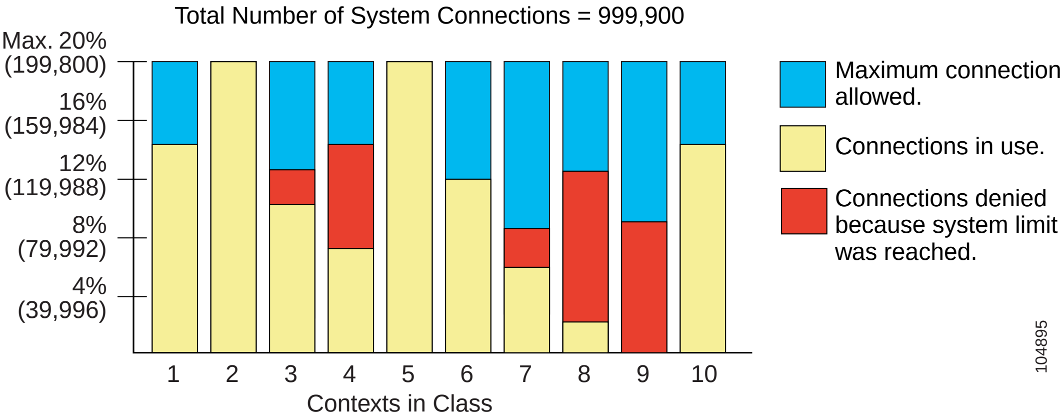

If you assign an absolute value to a resource across all contexts that exceeds the practical limit of the security appliance, then the performance of the security appliance might be impaired.

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 11-7.) Setting unlimited access is similar to oversubscribing the security appliance, except that you have less control over how much you oversubscribe the system.

Figure 11-7 Unlimited Resources

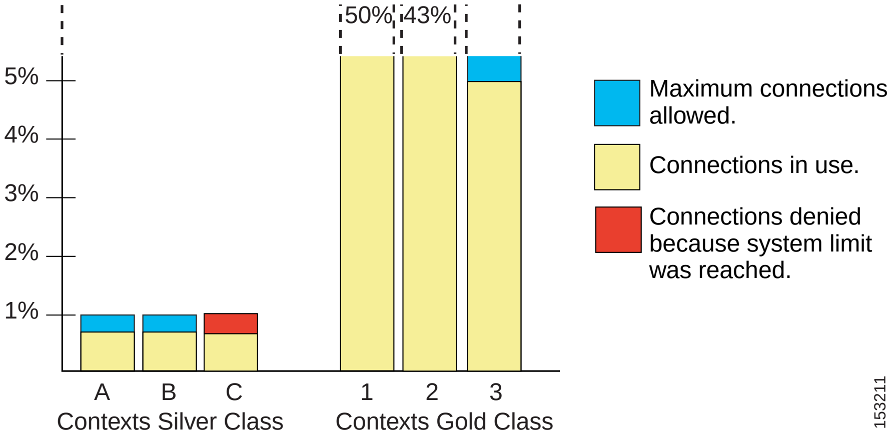

Default Class

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 11-8 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 11-8 Resource Classes

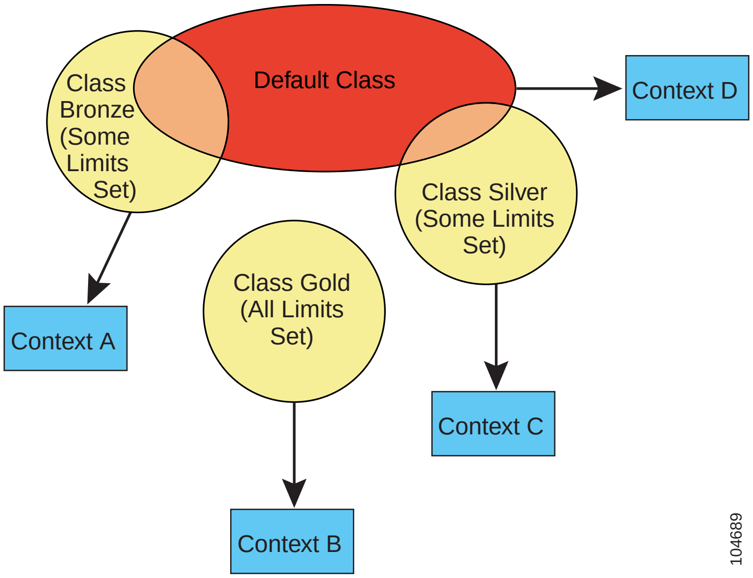
Class Bronze (Some Limits Set)
Default Class
Context D
Class Silver (Some Limits Set)
Class Gold (All Limits Set)
Context A
Context C
Context B
104689

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

For more information about resource classes, see the "Classes and Class Members Overview" section on page 11-11.

To add a resource class, perform the following steps:

Step 1 If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.

Step 2 On the Context Management > Resource Class pane, click **Add**.

The Add Resource Class dialog box appears.

Step 3 In the Resource Class field, enter a class name up to 20 characters in length.

Step 4 In the Count Limited Resources area, set 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.

You can set one or more of the following limits:

- 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 *Release Notes for Cisco ASDM* 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.

Step 5 In the Rate Limited Resources area, set 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.

You can set one or more of the following limits:

- 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.

Step 6 Click **OK**.

Monitoring Context Resource Usage

To monitor resource usage of all contexts from the system execution space, perform the following steps:

Step 1 If you are not already in the System mode, in the Device List pane, double-click **System** under the active device IP address.

Step 2 Click the **Monitoring** button on the toolbar.

Step 3 Click **Context Resource Usage**.

Click each resource type to view the resource usage for all contexts:

- **ASDM**—Shows the usage of ASDM connections.
  - Context—Shows the name of each context.
  - Existing Connections (#)—Shows the number of existing connections.
  - Existing Connections (%)—Shows the connections used by this context as a percentage of the total number of connections used by all contexts.
  - Peak Connections (#)—Shows the peak number of connections since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.
- **Telnet**—Shows the usage of Telnet connections.
  - Context—Shows the name of each context.
  - Existing Connections (#)—Shows the number of existing connections.
  - Existing Connections (%)—Shows the connections used by this context as a percentage of the total number of connections used by all contexts.
  - Peak Connections (#)—Shows the peak number of connections since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.
- **SSH**—Shows the usage of SSH connections.
  - Context—Shows the name of each context.
  - Existing Connections (#)—Shows the number of existing connections.
  - Existing Connections (%)—Shows the connections used by this context as a percentage of the total number of connections used by all contexts.
  - Peak Connections (#)—Shows the peak number of connections since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.
- **Xlates**—Shows the usage of netword address translations.
  - Context—Shows the name of each context.
  - Xlates (#)—Shows the number of current xlates.
  - Xlates (%)—Shows the xlates used by this context as a percentage of the total number of xlates used by all contexts.

  - Peak (#)—Shows the peak number of xlates since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.
- **NATs—**Shows the number of NAT rules.
  - Context—Shows the name of each context.
  - NATs (#)—Shows the current number of NAT rules.
  - NATs (%)—Shows the NAT rules used by this context as a percentage of the total number of NAT rules used by all contexts.
  - Peak NATs (#)—Shows the peak number of NAT rules since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.
- **Syslogs**—Shows the rate of system log messages.
  - Context—Shows the name of each context.
  - Syslog Rate (#/sec)—Shows the current rate of system log messages.
  - Syslog Rate (%)—Shows the system log messages generated by this context as a percentage of the total number of system log messages generated by all contexts.
  - Peak Syslog Rate (#/sec)—Shows the peak rate of system log messages since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.

Step 4 Click **Refresh** to refresh the view.

Configuring Security Contexts

This section describes how to add security contexts, and includes the following topics:

- Adding a Security Context, page 11-16
- Automatically Assigning MAC Addresses, page 11-18

For more information about security contexts, see the "Security Context Overview" section on page 11-1.

Adding a Security Context

For more information about security contexts, see the "Security Context Overview" section on page 11-1.

To add a security context, perform the following steps:

Step 1 If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.

Step 2 On the Context Management > Security Contexts pane, click **Add**.

The Add Context dialog box appears.

Step 3 In the Security Context field, enter 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.

Step 4 In the Interface Allocation area, click the **Add** button to assign an interface to the context.

Step 5 From the Interfaces > Physical Interface drop-down list, choose an interface.

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.

Step 6 (Optional) In the Interfaces > Subinterface Range (optional) drop-down list, choose a subinterface ID.

For a range of subinterface IDs, choose the ending ID in the second drop-down list, if available.

In transparent firewall mode, only subinterfaces that have not been allocated to other contexts are shown.

Step 7 (Optional) In the Aliased Names area, check **Use Aliased Name in Context** to set an aliased name for this interface to be used in the context configuration instead of the interface ID.

a. In the Name field, sets the aliased name.

An aliased name must start with a letter, end with a letter, 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.

b. (Optional) In the Range field, set 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.

Step 8 (Optional) To enable context users to see physical interface properties even if you set an aliased name, check **Show Hardware Properties in Context**.

Step 9 Click **OK** to return to the Add Context dialog box.

Step 10 (Optional) If you use IPS virtual sensors, then assign a sensor to the context in the IPS Sensor Allocation area.

For detailed information about IPS and virtual sensors, see Chapter 30, "Configuring the IPS Application on the AIP SSM and SSC."

Step 11 (Optional) To assign this context to a resource class, choose a class name from the Resource Assignment > Resource Class drop-down list.

You can add or edit a resource class directly from this area. See the "Configuring Resource Classes" section on page 11-10 for more information.

Step 12 To set the context configuration location, identify the URL by choosing a file system type from the Config URL drop-down list and entering a path in the field.

For example, the combined URL for FTP has the following format:

ftp://server.example.com/configs/admin.cfg

Step 13 (Optional) For external filesystems, set the username and password by clicking **Login**.

Step 14 (Optional) To set the failover group for active/active failover, choose the group name in the Failover Group drop-down list.

Step 15 (Optional) Add a description in the Description field.

Automatically Assigning MAC Addresses

This section tells how to configure auto-generation of MAC addresses, and includes the following sections:

- Information About MAC Addresses, page 11-18
- Default MAC Address, page 11-18
- Failover MAC Addresses, page 11-18
- MAC Address Format, page 11-19
- Enabling Auto-Generation of MAC Addresses, page 11-19
- Viewing Assigned MAC Addresses, page 11-19

Information About MAC Addresses

To allow contexts to share interfaces, we suggest that you assign unique MAC addresses to each shared 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 11-2 for information about classifying packets.

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 Chapter 9, "Configuring Interfaces," to manually set the MAC address.

Default MAC Address

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.

All auto-generated MAC addresses start with A2. The auto-generated MAC addresses are persistent across reloads.

Interaction with Manual MAC Addresses

If you manually assign a MAC address and also enable auto-generation, then the manually assigned MAC address is used. If you later remove the manual MAC address, the auto-generated address is used.

Because auto-generated addresses start with A2, you cannot start manual MAC addresses with A2 if you also want to use auto-generation.

Failover MAC Addresses

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. See the "MAC Address Format" section for more information.

For upgrading failover units with the legacy version of the **mac-address auto** command before the **prefix** keyword was introduced, see the **mac-address auto** command in the *Cisco ASA 5500 Series Command Reference*.

MAC Address Format

The security appliance generates the MAC address using the following format:

A2xx.yyzz.zzzz

Where xx.yy is a user-defined prefix, and zz.zzzz is an internal counter generated by the security appliance. For the standby MAC address, the address is identical except that the internal counter is increased by 1.

For an example of how the prefix is used, if you set a prefix of 77, then the security appliance converts 77 into the hexadecimal value 004D (yyxx). When used in the MAC address, the prefix is reversed (xxyy) to match the security appliance native form:

A2**4D.00**zz.zzzz

For a prefix of 1009 (03F1), the MAC address is:

A2**F1.03**zz.zzzz

Enabling Auto-Generation of MAC Addresses

To enable automatic MAC address assignment, perform the following steps:

Step 1 If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.

Step 2 Choose the **Configuration** > **Context Management** > **Security Contexts** pane, and check **Mac-Address auto**.

Step 3 Check the **Prefix** check box, and in the field, enter a a decimal value between 0 and 65535.

This prefix is converted to a 4-digit hexadecimal number, and used as part of the MAC address. The prefix ensures that each security appliance uses unique MAC addresses, so you can have multiple security appliances on a network segment, for example. See the "MAC Address Format" section for more information about how the prefix is used.

When you configure a name for the interface in a context, the new MAC address is generated immediately. If you enable this option after you configure context interfaces, then MAC addresses are generated for all interfaces immediately after you apply it. If you disable the 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.

Note For the MAC address generation method when not using a prefix (not recommended), see the **mac-address auto** command in the *Cisco ASA 5500 Series Command Reference*.

Viewing Assigned MAC Addresses

You can view auto-generated MAC addresses within the system configuration or within the context. This section includes the following topics:

- Viewing MAC Addresses in the System Configuration, page 11-20
- Viewing MAC Addresses Within a Context, page 11-20

Viewing MAC Addresses in the System Configuration

To view the assigned MAC addresses from the system execution space, perform the following steps:

Step 1 If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.

Step 2 Choose the **Configuration > Context Management > Security Contexts** pane, and view the Primary MAC and Secondary MAC columns.

Note If you manually assign a MAC address to an interface, but also have auto-generation enabled, the auto-generated address continues to show in the configuration even though the manual MAC address is the one that is in use. If you later remove the manual MAC address, the auto-generated one shown will be used.

Viewing MAC Addresses Within a Context

To view the MAC address in use by each interface within the context,

Step 1 If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.

Step 2 Choose the **Configuration > Interfaces** pane, and view the MAC Address address column.

This table shows the MAC address in use; if you manually assign a MAC address and also have auto-generation enabled, then you can only view the unused auto-generated address from within the system configuration.

CHAPTER 12

Configuring Dynamic And Static Routing

To configure static routes and dynamic routing protocols, go to **Configuration** > **Device Setup** > **Routing** area of the ASDM interface.

You can configure up to two OSPF, one EIGRP, and one RIP routing process on the security appliance at the same time. Dynamic routing is only available on security appliances in routed firewall mode; you cannot configure dynamic routing protocols on a security appliance in transparent firewall mode.

You can configure static routes on security appliances in either routed or transparent firewall mode. You can use the static route tracking feature to have the security appliance a backup static route if a primary static route becomes unavailable.

This section contains the following topics:

- Dynamic Routing, page 12-1
- Static Routes, page 12-40
- ASR Group, page 12-46
- Proxy ARPs, page 12-47

Dynamic Routing

This section contains the following topics:

- OSPF, page 12-1
- RIP, page 12-22
- EIGRP, page 12-28

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 instead of the 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, one RIP routing process, and one EIGRP routing process running on the security appliance at the same time.

For more information about enabling and configuring OSPF, see the following:

- Setup, page 12-2
- Filtering, page 12-8
- Interface, page 12-10
- Redistribution, page 12-15
- Static Neighbor, page 12-17
- Summary Address, page 12-18
- Virtual Link, page 12-20

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, page 12-3
- Setup > Area/Networks Tab, page 12-5
- Setup > Route Summarization Tab, page 12-7

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

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. 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. See Edit OSPF Process Advanced Properties, page 12-3 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
•	—	•	—	—

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

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

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

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

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

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

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

The Interface pane lets you configure interface-specific OSPF routing properties, such as OSPF message authentication and 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

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

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

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

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

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

The Redistribution pane displays the rules for redistributing routes from one routing process into an OSPF routing process.

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—Static routes are redistributed into the OSPF routing process.
  - Connected—The route was established automatically by virtue of having IP enabled on the interface. These routes are redistributed into the OSPF routing process as external to the AS.
  - OSPF—Routes from another OSPF routing process are being redistributed into the OSPF routing process.
  - EIGRP—Routes are redistributed from the EIGRP routing process into the OSPF routing process.
  - RIP—Routes are redistributed from the RIP routing process into the OSPF routing process.
- Match—Displays the conditions used for redistributing routes from one OSPF routing process 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

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—Redistribute static routes into the OSPF routing process.
  - Connected—Redistribute connected routes (routes established automatically by virtue of having IP enabled on the interface) into the OSPF routing process. Connected routes are redistributed as external to the AS.
  - OSPF—Redistribute routes from another OSPF routing process. Choose the OSPF process ID from the list.
  - RIP—Redistribute routes from the RIP routing process.
  - EIGRP—Redistribute routes from the EIGRP routing process. Choose the autonomous system number of the EIGRP routing process from the list.
- Match—Displays the conditions used for redistributing routes from another OSPF routing process into the selected OSPF routing process. These options are not available when redistributing static, connected, RIP, or EIGRP routes. 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—Check 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

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

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, page 12-40.)

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

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

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 checked.
- 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

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

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

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

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.

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.

Setup

Use the 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 version of RIP used by the security appliance. If this check box is 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

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—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

The Edit RIP Interface Entry dialog box allows you to configure the 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 unsupported 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 contain 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

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—Clicking 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

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

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
•	—	•	—	—

Redistribution

The 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.
  - EIGRP—Networks discovered by the specified EIGRP 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

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.
  - EIGRP and EIGRP ID—Routes discovered by the EIGRP routing process. If you choose EIGRP, you must also specify the autonomous system number of the EIGRP routing process in the EIGRP ID field.
- 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 check box to specify a metric for the redistributed routes. If not specified, the routes are assigned a metric of 0.
  - Transparent—Choose this option to cause the current route metric to be used.
  - 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 redistributed.
  - 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
•	—	•	—	—

EIGRP

EIGRP is an enhanced version of IGRP developed by Cisco. Unlike IGRP and RIP, EIGRP does not send out periodic route updates. EIGRP updates are sent out only when the network topology changes.

You can enable only one EIGRP routing process on the security appliance.

This section contains the following information:

- Configuring EIGRP, page 12-29
- Field Information for the EIGRP Panes, page 12-30

For information about monitoring dynamically discovered EIGRP neighbors, see Monitoring EIGRP Neighbors, page 44-7.

Configuring EIGRP

To configure EIGRP routing on the Security Appliance, perform the following steps:

Step 1 Go to the **Configuration > Device Setup > Routing > EIGRP** area of the ASDM interface.

Step 2 Enable the EIGRP routing process on the **Setup > Process Instances** tab. See Process Instances, page 12-30 for more information.

Step 3 (Optional) Configure the EIGRP routing process parameters. Click **Advanced** on the Setup > Process Instances tab.

You can configure the EIGRP routing process as a stub routing process, disable automatic route summarization, define the default metrics for redistributed routes, change the administrative distances for internal and external EIGRP routes, configure a static router ID, and enable or disable the logging of adjacency changes. See Edit EIGRP Process Advanced Properties, page 12-31 for more information.

Step 4 Define the networks and interfaces that will participate in EIGRP routing on the **Setup > Networks** tab. See Networks, page 12-32 for more information.

Directly-connected and static networks that fall within the defined networks are advertised by the security appliance. Additionally, only interfaces with an IP address that fall within the defined networks participate in the EIGRP routing process.

If you have an interface that you do not want to participate in EIGRP routing, but that is attached to a network that you want advertised, configure a network entry on the Setup > Networks tab that covers the network the interface is attached to, and then configure that interface as a passive interface to prevent the interface from sending or receiving EIGRP updates. Interfaces configured as passive do not send or receive EIGRP updates. See Passive Interfaces, page 12-33 for more information.

Step 5 (Optional) Define route filters on the **Filter Rules** pane. Route filtering provides more control over the routes that are allowed to be sent or received in EIGRP updates. See Filter Rules, page 12-34 for more information.

Step 6 (Optional) Define route redistribution on the **Redistribution** pane.

You can redistribute routes discovered by RIP and OSPF into the EIGRP routing process. You can also redistribute static and connected routes into the EIGRP routing process. When redistributing static and connected routes into the EIGRP routing process, metrics are not required to be configured, although this is recommended.You do not need to redistribute static or connected routes if they fall within the range of a network configured on the Setup > Networks tab. See Redistribution, page 12-36 for more information.

Step 7 (Optional) Define static EIGRP neighbors on the **Static Neighbor** pane.

EIGRP hello packets are sent as multicast packets. If an EIGRP neighbor is located across a nonbroadcast network, such as a tunnel, you must manually define that neighbor. When you manually define an EIGRP neighbor, hello packets are sent to that neighbor as unicast messages. See Static Neighbor, page 12-37 for more information.

Step 8 (Optional) Define summary addresses on the **Summary Address** pane.

You need to manually define summary addresses if you want to create summary addresses that do not occur at a network number boundary or if you want to use summary addresses on a security appliance with automatic route summarization disabled. See Summary Address, page 12-38 for more information about defining summary addresses. See Edit EIGRP Process Advanced Properties, page 12-31 for information about enabling and disabling automatic route summarization.

Step 9 (Optional) Define interface-specific EIGRP parameters on the **Interfaces** pane. These parameters include EIGRP message authentication, hold time, hello interval, delay metric, and the use of split-horizon. See Interface, page 12-35 for more information.

Step 10 (Optional) Control the sending and receiving of default route information in EIGRP updates on the **Default Information** pane. By default, default routes are sent and accepted. See Default Information, page 12-39 for more information.

Field Information for the EIGRP Panes

This section contains the following topics:

- Setup, page 12-30
- Filter Rules, page 12-34
- Interface, page 12-35
- Redistribution, page 12-36
- Static Neighbor, page 12-37
- Summary Address, page 12-38
- Default Information, page 12-39

Setup

Enable and EIGRP process and configure the basic setting for that process on the Setup pane. The Setup pane contains the following tabs:

- Process Instances, page 12-30
- Networks, page 12-32
- Passive Interfaces, page 12-33

Modes

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

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

Process Instances

The Process Instances tab lets you enable an EIGRP routing process.

Fields

- Enable this EIGRP Process—Check this check box to enable an EIGRP routing process. You can only enable one EIGRP routing process on the device. You must enter an autonomous system number for the routing process in the EIGRP Process field before you can save your change.
- EIGRP Process—Enter the autonomous system number for the EIGRP process. The autonomous system number can be from 1 to 65535.

- Advanced—Click this button to configure the EIGRP process settings, such as the router ID, default metrics, stub routing settings, neighbor change and warning logging, and the administrative distances for the EIGRP routes.

Modes

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

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

Edit EIGRP Process Advanced Properties

The Edit EIGRP Process Advanced Properties dialog box lets you configure the router ID, default metrics, stub routing settings, neighbor change and warning logging, and the administrative distances of the EIGRP routes for the EIGRP routing process.

Fields

- EIGRP—*Display only.* Displays the autonomous system number for the EIGRP routing process.
- Router Id—Enter an IP address to be used as the Router ID for the security appliance in the EIGRP routing process. The router ID is used to identify the originating router for external routes. The IP address does not have to an address configured on the security appliance, however it must be unique within the routing domain. If not specified, the highest-level IP address on the security appliance is used as the router ID.
- Auto-Summary—Check to enable automatic route summarization. Clear to disable automatic route summarization. This setting is enabled by default.
- Default Metrics—The default metrics are applied to routes redistributed into the EIGRP routing process. If not specified, you must specify the metrics when you configure the redistribution (see Redistribution, page 12-36).
  - Bandwidth—The minimum bandwidth of the route in kilobytes per second. Valid values are from 1 to 4294967295.
  - Loading—The effective bandwidth of the route expressed as a number from 1 to 255 (255 is 100 percent loading).
  - Reliability—The likelihood of successful packet transmission expressed as a number from 0 through 255. The value 255 means 100 percent reliability; 0 means no reliability.
  - Delay—The route delay in tens of microseconds. Valid values are 1 to 4294967295.
  - MTU—The smallest allowed value for the maximum transmission unit, expressed in bytes. Valid values are from 1 to 65535.
- Stub—The stub area contains the setting for creating an EIGRP stub routing process. A stub routing process does not maintain a full topology table. At a minimum, stub routing needs a default route to a distribution router, which makes the routing decisions.
  - Stub Receive only—Configures the EIGRP stub routing process to receive route information from the neighbor routers but does not send route information to the neighbors. If this option is selected, you cannot select any of the other stub routing options.

  - Stub Connected—Advertises connected routes.
  - Stub Static—Advertises static routes.
  - Stub Redistributed—Advertises redistributed routes.
  - Stub Summary—Advertises summary routes.
- Adjacency Changes—Lets you configure the logging of neighbor warning and change messages. Logging for both is enabled by default.
  - Log Neighbor Changes—Check to enable or uncheck to disable the logging of neighbor adjacency changes.
  - Log Neighbor Warnings—Check to enable or uncheck to disable the logging of neighbor adjacency changes. Enter the time interval (in seconds) between repeated neighbor warning messages. Valid values are from 1 to 65535. Repeated warnings are not logged if they occur during this interval.
- Administrative Distance—Lets you configure the administrative distances for internal and external EIGRP routes.
  - Internal Distance—Administrative distance for EIGRP internal routes. Internal routes are those that are learned from another entity within the same autonomous system. Valid values are from 1 to 255. The default value is 90.
  - External Distance—Administrative distance for EIGRP external routes. External routes are those for which the best path is learned from a neighbor external to the autonomous system. Valid values are from 1 to 255. The default value is 170.

Modes

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

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

Networks

The Network tab lets you specify the networks used by the EIGRP routing process. For an interface to participate in EIGRP routing, it must fall within the range of addresses defined by the network entries. For directly connected and static networks to be advertised, they must also fall within the range of the network entries.

The Network table displays the networks configure for the EIGRP routing process. Each row of the table displays the network address and associated mask configure for the specified EIGRP routing process. To add or change a network, do one of the following:

- To add a new network entry, click **Add**. The Add EIGRP Network dialog box appears.
- To remove a network entry, select the entry in the table and click **Delete**.
- To change a network entry, you must first remove the entry and then add a new one. You cannot edit existing entries.

Fields

The Add EIGRP Network Entry dialog box fields:

- EIGRP AS—Displays the autonomous system number of the EIGRP routing process.
- IP Address—Enter the IP address of the networks to participate in the EIGRP routing process.
- Network Mask—Select or enter a network mask to apply to the 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
•	—	•	—	—

Passive Interfaces

The Passive Interface tab lets you configure one or more interfaces as passive interfaces. In EIGRP, a passive interface does not send or receive routing updates.

The Passive Interface table lists each interface configured as a passive interface. To configure whether an interface participates in EIGRP routing, do one of the following:

- To specify all interfaces as passive, check the Suppress routing updates on all interfaces check box. Even if an interface is not shown in the Passive Interface table, it will be configured as passive when this check box is selected.
- To add a passive interface entry, click **Add**. The Add EIGRP Passive Interface dialog box appears. You can select the interface you want to make passive in the dialog box.
- To remove a passive interface, select the interface in the table and click **Delete**.

Fields

Passive Interface pane fields:

- EIGRP Process—The autonomous system number of the EIGRP routing process.
- Suppress routing updates on all interfaces—Check this check box to set all interfaces to passive. Clear this check box to allow all interfaces to send and receive EIGRP updates. Note that the interfaces must also have an associated network entry to participate in EIGRP routing.
- Passive Interfaces table—Displays the interface configured as passive.
  - Interface—Displays the name of the interface.
  - EIGRP Process—Displays the autonomous system number of the EIGRP process.
  - Passive—Displays “true” to indicate that the interface is operating in passive mode.

Add Passive Interface dialog box fields:

- EIGRP AS—The autonomous system number of the EIGRP routing process.
- Interface—Select the interface 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
•	—	•	—	—

For More Information

- Configuring EIGRP, page 12-29

Filter Rules

The Filter Rules pane displays the route filtering rules configured for the EIGRP routing process. Filter rules let you control which routes are accepted or advertised by the EIGRP routing process.

Each row of the Filter Rule table describes a filter rule for a specific interface or routing protocol. For example, a filter rule with a direction of “in” on the outside interface would apply filtering to any EIGRP updates received on the outside interface. A filter rule with a direction of “out” with OSPF 10 specified as the routing protocol would apply the filter rules to routes redistributed into the EIGRP routing process in outbound EIGRP updates.

To configure filter rules, do one of the following:

- To add a filter rule, click **Add**. The Add Filter Rules dialog box appears.
- To edit a filter rule, select the filter rule in the table and click **Edit**. You can also double-click a filter rule to edit the rule. The Edit Filter Rules dialog box appears.
- To remove a filter rule, select the filter rule in the table and click **Delete**.

Fields

The Add/Edit EIGRP Filter Rule Dialog box fields:

- EIGRP—The autonomous system number of the EIGRP routing process.
- Direction—Select “in” for rules that filter routes from incoming EIGRP routing updates. Select “out” to filter routes from EIGRP routing updates sent by the security appliance.
- Routing process—(For outgoing filters only) Specifies the type of route being filtered. You can filter routes redistributed from static, connected, RIP, and OSPF routing processes. Filters that specify a routing process filter those routes from updates sent on all interfaces.
- Id—The OSPF process ID.
- Interface—The interface the filter applies to.
- Add—Opens the Network Rule dialog box.
- Edit—Opens the Network Rule dialog box for the selected network rule.

Add/Edit Network Rule dialog box lets you define an access list for the filter rule. The dialog box contains the following fields:

- Action—Select Permit to allow the specified network to be advertised. Select Deny to prevent the specified network from being advertised.
- IP Address—Type IP address of the network being permitted or denied. To permit or deny all addresses, use the IP address 0.0.0.0 with a network mask of 0.0.0.0.
- 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
•	—	•	—	—

For More Information

- Configuring EIGRP, page 12-29

Interface

The Interface pane displays the EIGRP interface configurations. The Interface Parameters table displays all of the interfaces on the security appliance and lets you modify the following settings on a per-interface basis:

- Authentication key and mode.
- The EIGRP hello interval and hold time.
- The interface delay metric used in EIGRP metric calculations.
- The use of split-horizon on the interface.

To configure the EIGRP parameters for an interface, double-click an interface entry or select the entry and click **Edit**. The Edit EIGRP Interface Entry dialog box appears.

Fields

The Edit EIGRP Interface Entry dialog box fields:

- Interface—*Display only.* Displays the interface being modified.
- AS—The EIGRP autonomous system number.
- Hello Interval—Enter the interval between EIGRP hello packets sent on an interface. Valid values are from 1 to 65535 seconds. The default value is 5 seconds.
- Hold Time—Specifies the hold time, in seconds. Valid values are from 1 to 65535 seconds. The default value is 15 seconds.
- Split Horizon—Check this check box to enable split horizon on the interface. Uncheck the check box to disable split horizon. Split horizon is enabled by default.
- Delay—Enter the delay value in this field. The delay time is in tens of microseconds. Valid values are from 1 to 16777215.
- Enable MD5 Authentication—Check this check box to enable MD5 authentication of EIGRP process messages.
  - Key—Key to authenticate EIGRP updates. The key can contain up to 16 characters.
  - Key ID—Key identification value; valid values range from 1 to 255.

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 EIGRP, page 12-29

Redistribution

The Redistribution pane displays the rules for redistributing routes from other routing protocols into the EIGRP routing process. When redistributing static and connected routes into the EIGRP routing process, metrics are not required to be configured, although this is recommended.

Each row of the Redistribution pane table contains a route redistribution entry.

To add or modify route redistribution into the EIGRP routing process, do one of the following:

- To add a new redistribution rule, click **Add**. The Add EIGRP Redistribution Entry dialog box opens.
- To edit an existing EIGRP static neighbor, select the address in the table and click **Edit**. You can also double-click an entry in the table to edit that entry. The Edit EIGRP Redistribution Entry dialog box opens.

Fields

The Add/Edit EIGRP Redistribution Entry dialog box fields:

- AS—Displays the autonomous system number of the EIGRP routing process to which the entry applies.
- Static—Redistributes static routes into the EIGRP routing process. Static routes that fall within the scope of a network statement are automatically redistributed into EIGRP; you do not need to define a redistribution rule for them.
- Connected—Redistributes connected routes into the EIGRP routing process. Connected routes that fall within the scope of a network statement are automatically redistributed into EIGRP; you do not need to define a redistribution rule for them.
- RIP—Redistributes routes discovered by the RIP routing process into EIGRP.
- Optional Metrics—Defines the metrics used for the redistributed route. You do not need to define these values if you already defined the default metrics in the **Edit EIGRP Process Advanced Properties** dialog box (see Edit EIGRP Process Advanced Properties, page 12-31 for information about setting the default metrics).
  - Bandwidth—EIGRP bandwidth metric in Kilobits per second. Valid values are from 1 to 4294967295.
  - Delay—EIGRP delay metric, in 10 microsecond units. Valid values are from 0 to 4294967295.
  - Reliability—EIGRP reliability metric. Valid values are from 0 to 255, where 255 indicates 100% reliability.
  - Loading—EIGRP effective bandwidth (loading) metric. Valid values are from 1 to 255, where 255 indicates 100% loaded.
  - MTU—The MTU of the path. Valid values are from 1 to 65535.

- Route Map—To further define which routes are redistributed into the EIGRP routing process, enter the name of a route map.
- Optional OSPF Redistribution—these options let you further specify which OSPF routes are redistributed into the EIGRP routing process.
  - Match Internal—Match routes internal to the specified OSPF process.
  - Match External 1—Match type 1 routes external to the specified OSPF process.
  - Match External 2—Match type 2 routes external to the specified OSPF process.
  - Match NSSA-External 1—Match type 1 routes external to the specified OSPF NSSA.
  - Match NSSA-External 2—Match type 2 routes external to the specified OSPF NSSA.

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 EIGRP, page 12-29

Static Neighbor

The Static Neighbor pane displays the statically-defined EIGRP neighbors. An EIGRP neighbor sends EIGRP routing information to and receives EIGRP routing information from the security appliance. Normally, neighbors are dynamically discovered through the neighbor discovery process. However, on point-to-point, non-broadcast networks, you must statically define the neighbors.

Each row of the Static Neighbor table displays the EIGRP autonomous system number for the neighbor, the neighbor IP address, and the interface through which the neighbor is available.

To configure a static neighbor, so one of the following:

- To add a new EIGRP static neighbor, click **Add**. The Add EIGRP Neighbor Entry dialog box opens.
- To edit an existing EIGRP static neighbor, select the address in the table and click **Edit**. You can also double-click an entry in the table to edit that entry. The Edit EIGRP Neighbor Entry dialog box opens.

Fields

The Add/Edit EIGRP Neighbor Entry dialog box fields:

- EIGRP AS—The autonomous system number for the EIGRP process the neighbor is being configured for.
- Interface Name—Select the interface through which the neighbor is available from the list.
- Neighbor IP Address—Enter the IP address of the neighbor.

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 EIGRP, page 12-29

Summary Address

The Summary Address pane displays a table of the statically-defined EIGRP summary addresses. By default, EIGRP summarizes subnet routes to the network level. You can create statically-defined EIGRP summary addresses to the subnet level from the Summary Address pane.

To create or modify a summary address, do one of the following:

- To add a new EIGRP summary address, click **Add**. The Add Summary Address dialog box opens.
- To edit an existing EIGRP summary address, select the address in the table and click **Edit**. You can also double-click an entry in the table to edit that entry. The Edit Summary Address dialog box opens.

Fields

The Add/Edit EIGRP Summary Address Entry dialog box contains the following fields. These fields are also shown in the Summary Address table.

- EIGRP AS—Select the autonomous system number of the EIGRP routing process the summary address applies to.
- Interface—The interface the summary address is advertised from.
- IP Address—Enter the IP address of the summary route.
- Netmask—Select or enter the network mask to apply to the IP address.
- Administrative Distance—Enter the administrative distance for the route. If left blank, the route has the default administrative distance of 5.

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 EIGRP, page 12-29

Default Information

The Default Information pane displays a table of rules for controlling the sending and receiving of default route information in EIGRP updates. You can have one "in" and one "out" rule for each EIGRP routing process (only one process is currently supported).

By default, default routes are sent and accepted. To restrict or disable the sending and receiving of default route information, perform the following steps:

Step 1 Open the **Configuration > Device Setup > Routing > EIGRP > Default Information** pane.

Step 2 Do one of the following:

- To create a new entry, click **Add**.
- To edit an entry, double-click the entry in the table or select an entry in the table an click **Edit**.

The Add or Edit Default Information dialog box opens for that entry. The EIGRP autonomous system number is automatically selected in the EIGRP field.

Step 3 Set the direction for the rule in the **Direction** field:

- in—the rule filters default route information from incoming EIGRP updates.
- out—the rule filters default route information from outgoing EIGRP updates.

You can have one "in" rule and one "out" rule for each EIGRP process.

Step 4 Add network rules to the network rule table. The network rules define which networks are allowed and which are not when receiving or sending default route information. Repeat the following steps for each network rule you are adding to the default information filter rule.

a. Click **Add** to add a network rule. Double-click an existing network rule to edit the rule.

b. In the **Action** field, select **Permit** to allow the network or **Deny** to block the network.

c. Enter the IP address and network mask of the network being permitted or denied by the rule in the **IP Address** and **Network Mask** fields.

   To deny all default route information from being accepted or sent, use 0.0.0.0 as the network address and select 0.0.0.0 as the network mask.

d. Click **OK** to add the specified network rule to the default information filter rule.

Step 5 Click **OK** to accept the default information filter rule.

Fields

Add/Edit Default Information dialog box:

- EIGRP—Select the autonomous system number of the EIGRP routing process the default information filter applies to.
- Direction—Select "in" to filter default route information from incoming route updates. Select "out" to filter default route information from outgoing route updates.
- Add—Add a network rule to the default information filter rule.
- Edit—Modify an existing network rule.

Network Rule dialog box. The network rules appear in the Filter Rules column of the Default Information filter rule table.

- Action—Select Permit to allow the specified network to be advertised. Select Deny to prevent the specified network from being advertised.

- IP Address—Type IP address of the network being permitted or denied. To permit or deny all addresses, use the IP address 0.0.0.0 with a network mask of 0.0.0.0.
- 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
•	—	•	—	—

For More Information

- Configuring EIGRP, page 12-29

Static Routes

This section includes the following topics:

- Information About Static Routes, page 12-40
- Configuring Static Routes, page 12-41
- Information About Static Route Tracking, page 12-44
- Configuring Static Route Tracking, page 12-44
- Editing Static Routes, page 12-45
- Deleting Static Routes, page 12-46

Information About 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 that the security appliance knows out of which interface to send traffic. Traffic that originates on the security appliance might include communication 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 Information About Static Route Tracking, page 12-44, for the exception to this rule.) 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 for each interface. Equal Cost Multiple Path (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, you will 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 the "Configuring Static Routes" section on page 12-41.

Configuring Static Routes

You can create static routes that will access networks connected to a router on any interface. This procedure provides steps for defining static routes for any networks to which the security appliance is not directly connected. For information about static routes, see the "Information About Static Routes" section on page 12-40.

To configure a static route, perform the following steps:

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **Static Routes**.

Step 2 Choose which route to filter by clicking the appropriate radio button:

- **Both** (filters both IPv4 and IPv6)
- **IPv4 only**
- **IPv6 only**

Step 3 Click **Add**.

The Add Static Route window appears.

Step 4 From the Interface drop down list, choose the internal or external network interface name enabled in Interfaces:

- **management** (internal interface)
- **outside** (external interface)

Step 5 Enter an internal or external network IP address.

For IPv4 addresses, enter **0.0.0.0** to specify a default route. The 0.0.0.0 IP address can be abbreviated as 0. Optionally, click the ellipsis to browse for an address.

For IPv6 addresses, enter two colons (**::**) to specify a default route. Optionally, click the ellipsis to browse for an address.

Step 6 Enter the IP address of the gateway router, which is the next hop address for this route.

To enter a default route, set the IP address and mask to 0.0.0.0, or the shortened form of 0.

Optionally, click the ellipsis to browse for an address.

Note 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.

Step 7 Depending upon which route you chose to filter (IPv4, IPv6, or both), do one of the following:

- For IPv4 static routes (or for both IPv4 and IPv6 static routes), enter 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**.
- For IPv6 static routes only, enter a prefix length.

Step 8 Enter the administrative distance of the route.

The default is 1 if a metric is not specified. Leave the default setting of 1 unless you are sure of the number of hops to the gateway router.

Step 9 (Optional) In the Options field, select only one of these options for a static route.

- None—No options are specified for the static route. This option is the default.
- Tunneled—Specifies the route as the default tunnel gateway for VPN traffic. Use this option for the default route only. You can configure only 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 also display. The tracked option is supported in single, routed mode only.
  - Track ID—Enter 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—Enter a unique identifier for the SLA monitoring process.

Note The tracking option is not supported for IPv6.

Step 10 (Optional) Click the **Monitoring Options** button to display the Route Monitoring Options dialog box and change the following tracking object monitoring properties:

- 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, that 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.

Step 11 Click **OK**.

Step 12 Click **Apply** to save the configuration.

The monitoring process begins as soon as you save the newly configured route.

Modes

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

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

Modes

The following table shows the modes in which the Route Monitoring Options are available:

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

Information About 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.

Note Static route tracking is available for IPv4 routes only.

One problem with static routes is that there is no inherent mechanism for determining if the route is up or down. Static routes remain in the routing table even if the next hop gateway goes down. They are removed from the routing table only 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 feature 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 performs this action 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, 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 following targets:

- 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, with which the security appliance needs to communicate.
- 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 the "Configuring Static Route Tracking" section on page 12-44. To monitor the static route tracking process, see the "interface connection" section on page 42-9.

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 the "Configuring Static Routes" section on page 12-41.

Note Static route tracking is available for IPv4 routes only.

To configure tracking for a static route, perform the following steps:

Step 1 Choose a target of interest. Make sure the target responds to echo requests.

Step 2 Open the Static Routes page. Choose **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 the "Configuring Static Routes" section on page 12-41, for more information about these fields.

Step 4 Click the **Tracked** radio button 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 the "Configuring Static Routes" section on page 12-41, 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 Step 10 in the "Configuring Static Routes" section on page 12-41 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.

Editing Static Routes

To edit a static route, perform the following steps:

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **Static Routes**.

Step 2 From the main Static Routes pane, select which route to edit.

By default, the Both radio button is checked, and both IPv4 and IPv6 addresses appear in the pane.

- To limit your viewed choices to routes configured with IPv4 addresses, click the **IPv4** radio button.
- To limit your viewed choices to routes configured with IPv6 addresses, click the **IPv6** radio button.

Step 3 Click **Edit**.

The Edit Static Route window appears.

Step 4 Make the necessary configuration changes.

For information about entering values in specific fields, see the "Configuring Static Routes" section on page 12-41.

Step 5 Click **OK**.

The edited route information appears in the main Static Routes pane.

Step 6 Click **Apply** to save the changes to your configuration.

Deleting Static Routes

To delete a static route, perform the following steps:

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **Static Routes**.

Step 2 From the main Static Routes pane, select which route to delete.

By default, the Both radio button is checked, and both IPv4 and IPv6 addresses appear in the pane.

- To limit your viewed choices to routes configured with IPv4 addresses, click the **IPv4** radio button.
- To limit your viewed choices to routes configured with IPv6 addresses, click the **IPv6** radio button.

Step 3 Click **Delete**.

Step 4 The deleted route is removed from list of routes on in the main Static Routes pane.

Step 5 Click **Apply** to save the changes to your configuration.

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.

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.

Prerequisites

You must enable Stateful Failover for session information to be passed from the standby failover group to the active failover group.

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

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
•	•	•	•	—

CHAPTER 13

Configuring Multicast Routing

Multicast routing is supported in single, routed mode only. This section contains the following topics:

- Multicast, page 13-1—enable or disable multicast routing on the security appliance.
- IGMP, page 13-2—configure IGMP on the security appliance.
- Multicast Route, page 13-7—define static multicast routes.
- MBoundary, page 13-9—configure boundaries for administratively-scoped multicast addresses.
- MForwarding, page 13-11—enable or disable multicast forwarding on a per-interface basis.
- PIM, page 13-11—configure PIM on the security appliance.

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.

Note Only the UDP transport layer is supported for multicast routing.

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
•	—	•	—	—

For More Information

Configuring Multicast Routing, page 13-1

IGMP, page 13-2

Multicast Route, page 13-7

MBoundary, page 13-9

MForwarding, page 13-11

PIM, page 13-11

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

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

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

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

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

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

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

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

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

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

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. You can specify the interface, or the RPF neighbor, but not both at the same time.
- 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

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 239.255.255.255.255.
- 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

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 13-1

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 13-13
- Bidirectional Neighbor Filter, page 13-14
- Rendezvous Points
- Route Tree
- Request Filter

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

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 13-14

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 13-15

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

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

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

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

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

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

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 14

Configuring IPv6 Neighbors

This chapter provides information about IPv6 neighbor discovery. It shows how to add an IPv6 neighbor and how to configure neighbor solicitation messages.

This chapter includes the following sections:

Information About IPv6 Neighbor Discovery

Nodes (hosts) use neighbor discovery to determine the link-layer addresses for neighbors known to reside on attached links and to quickly purge cashed values that become invalid. Hosts also use neighbor discovery to find neighboring routers that are willing to forward packets on their behalf. Finally, nodes use the protocol to actively keep track of which neighbors are reachable and which are not, and to detect changed link-layer addresses. When a router or the path to a router fails, a host actively searches for functioning alternates.

The neighbor discovery process uses IPv6 (ICMPv6) messages and solicited-node multicast addresses. Every IPv6 node is required to join the multicast groups corresponding to its unicast and any cast addresses.

Neighbor solicitation messages are sent on the local link when a node wants to determine the link-layer address of another node on the same local link. This function is similar to the ARP in IPv4, but it avoids broadcasts used in IPv4 ARP messages, where all nodes receive unnecessary broadcast requests that do not concern them. The source node takes the right-most 24 bits of the IPv6 address of the destination node and sends a neighbor solicitation message to the solicited-node multicast group address on the local link. The destination node will respond with its link-layer address. To send a neighbor solicitation message, the source node must first identify the IPv6 unicast address of the destination node using a naming service mechanism, such as DNS. A neighbor solicitation message is also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified.

The neighbor advertisement message is a response to the neighbor solicitation message. After receiving the neighbor solicitation message, the destination node replies by sending a neighbor advertisement message on the local link. After receiving the neighbor advertisement, the source node and the destination node can communicate. Neighbor advertisement messages are also sent when there is a change in the link-layer address of a node on the local link.

A node can be manually added to the neighbor discovery cache.

Adding an IPv6 Static Neighbor

Ensure that IPv6 is enabled on at least one interface before attempting to add a neighbor, or ASDM will return an error message indicating that the configuration failed. For information about configuring IPv6 on an interface, see the Chapter 9, "Configuring Interfaces.".

For information about IPv6 Neighbor Discovery, see the "Information About IPv6 Neighbor Discovery" section on page 14-1.

To add an IPv6 static neighbor, perform the following steps:

Step 1 Choose **Configuration** > **Device Management** > **Advanced** > **IPv6 Neighbor Discovery Cache**.

Step 2 Click **Add**.

The Add IPv6 Static Neighbor dialog box appears.

Step 3 From the Interface Name drop-down list, choose an interface on which to add the neighbor.

Step 4 In the IP Address field, enter the IPv6 address that corresponds to the local data-link address, or click the ellipsis (**...**) to browse for an address.

If an entry for the specified IPv6 address already exists in the neighbor discovery cache—learned through the IPv6 neighbor discovery process—the entry is automatically converted to a static entry.

Step 5 In the MAC address field, enter the local data-line (hardware) MAC address.

Step 6 Click **OK**.

Note Before you apply the changes and save the configuration, you can click **Reset** to cancel any changes and restore the original values.

Step 7 Click **Apply** to save the configuration.

Configuring IPv6 Neighbor Discovery

The IPv6 neighbor discovery process uses ICMPv6 messages and solicited-node multicast addresses to determine the link-layer address of a neighbor on the same network (local link), verify the reachability of a neighbor, and keep track of neighboring routers. For information about adding an IPv6 static neighbor, see the "Adding an IPv6 Static Neighbor" section on page 14-2. For information about IPv6 neighbor discovery, see the "Information About IPv6 Neighbor Discovery" section on page 14-1.

This section includes the following topics:

- Configuring Neighbor Solicitation Messages, page 14-3
- Configuring Router Advertisement Messages, page 14-5

Configuring Neighbor Solicitation Messages

Neighbor solicitation messages (ICMPv6 Type 135) are sent on the local link by nodes attempting to discover the link-layer addresses of other nodes on the local link. The neighbor solicitation message is sent to the solicited-node multicast address.The source address in the neighbor solicitation message is the IPv6 address of the node sending the neighbor solicitation message. The neighbor solicitation message also includes the link-layer address of the source node.

After receiving a neighbor solicitation message, the destination node replies by sending a neighbor advertisement message (ICPMv6 Type 136) on the local link. The source address in the neighbor advertisement message is the IPv6 address of the node sending the neighbor advertisement message; the destination address is the IPv6 address of the node that sent the neighbor solicitation message. The data portion of the neighbor advertisement message includes the link-layer address of the node sending the neighbor advertisement message.

After the source node receives the neighbor advertisement, the source node and destination node can communicate. Figure 14-1 shows the neighbor solicitation and response process.

Figure 14-1 IPv6 Neighbor Discovery—Neighbor Solicitation Message

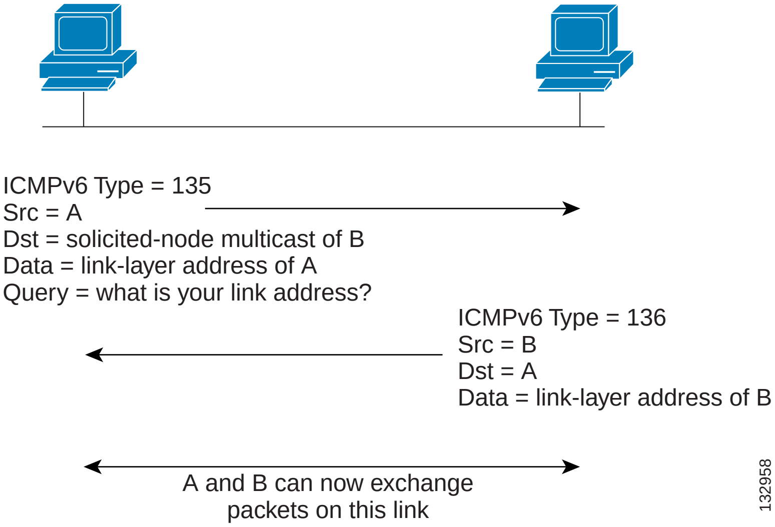
ICMPv6 Type = 135
Src = A
Dst = solicited-node multicast of B
Data = link-layer address of A
Query = what is your link address?
ICMPv6 Type = 136
Src = B
Dst = A
Data = link-layer address of B
A and B can now exchange
packets on this link
132958

Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. When a node wants to verifying the reachability of a neighbor, the destination address in a neighbor solicitation message is the unicast address of the neighbor.

Neighbor advertisement messages are also sent when there is a change in the link-layer address of a node on a local link. When there is such a change, the destination address for the neighbor advertisement is the all-nodes multicast address.

You can configure the neighbor solicitation message interval and neighbor reachable time on a per-interface basis. See the following topics for more information:

- Configuring the IPv6 Neighbor Solicitation Message Interval, page 14-4
- Configuring the IPv6 Neighbor Reachable Time, page 14-4

Configuring the IPv6 Neighbor Solicitation Message Interval

You can configure the interval between IPv6 neighbor solicitation retransmissions on an interface. Valid values range from 1000 to 3600000 milliseconds. The default value is 1000 milliseconds. This setting is also sent in router advertisement messages.

To configure the neighbor solicitation message interval, perform the following steps:

Step 1 Choose **Configuration** > **Device Setup** > **Interfaces**.

Step 2 Choose the interface on which to configure the neighbor solicitation interval. The interface must have been configured with an IPv6 address. See the "Configuring IPv6 Neighbor Discovery" section on page 14-2 for more information.

Step 3 Click **Edit**.

The Edit Interface dialog box appears with three tabs: General, Advanced, and IPv6.

Step 4 Click the **IPv6** tab.

Step 5 In the NS Interval field, enter the time interval.

Step 6 Click **OK**.

Step 7 Click **Apply** to save the configuration.

Configuring the IPv6 Neighbor Reachable Time

The neighbor reachable time enables detecting unavailable neighbors. Shorter configured times enable detecting unavailable neighbors more quickly; however, shorter times consume more IPv6 network bandwidth and processing resources in all IPv6 network devices. Very short configured times are not recommended in normal IPv6 operation.

Valid time values range from 0 to 3600000 milliseconds. The default is 0; however, when you use 0, the reachable time is sent as undetermined. It is up to the receiving devices to set and track the reachable time value.

To configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event has occurred, perform the following steps:

Step 1 Choose **Configuration** > **Device Setup** > **Interfaces**.

Step 2 Select the interface for which you want to configure the time. The interface must have been configured with an IPv6 address. See the "Configuring IPv6 Neighbor Discovery" section on page 14-2 for more information.

Step 3 Click **Edit**.

The Edit Interface dialog box appears with three tabs: General, Advanced, and IPv6.

Step 4 Click the **IPv6** tab.

Step 5 In the Reachable Time field, enter a valid value.

Step 6 Click **OK**.

Step 7 Click **Apply** to save the configuration.

Configuring Router Advertisement Messages

Router advertisement messages (ICMPv6 Type 134) are periodically sent out each IPv6 configured interface of the security appliance. The router advertisement messages are sent to the all-nodes multicast address.

Figure 14-2 IPv6 Neighbor Discovery—Router Advertisement Message

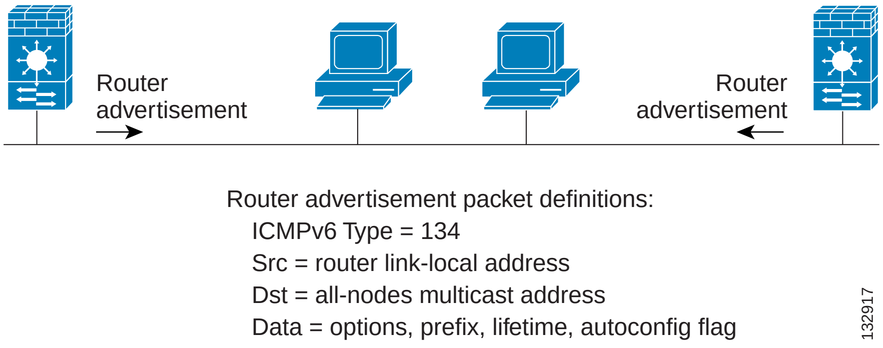

Router advertisement packet definitions:
ICMPv6 Type = 134
Src = router link-local address
Dst = all-nodes multicast address
Data = options, prefix, lifetime, autoconfig flag

Router advertisement messages typically include the following information:

- One or more IPv6 prefix that nodes on the local link can use to automatically configure their IPv6 addresses.
- Lifetime information for each prefix included in the advertisement.
- Sets of flags that indicate the type of autoconfiguration (stateless or stateful) that can be completed.
- Default router information (whether the router sending the advertisement should be used as a default router and, if so, the amount of time (in seconds) the router should be used as a default router).
- Additional information for hosts, such as the hop limit and MTU a host should use in packets that it originates.
- The amount of time between neighbor solicitation message retransmissions on a given link.
- The amount of time a node considers a neighbor reachable.

Router advertisements are also sent in response to router solicitation messages (ICMPv6 Type 133). Router solicitation messages are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled router advertisement message. Because router solicitation messages are usually sent by hosts at system startup, and the host does not have a configured unicast address, the source address in router solicitation messages is usually the unspecified IPv6 address (0:0:0:0:0:0:0:0). If the host has a configured unicast address, the unicast address of the interface sending the router solicitation message is used as the source address in the message. The destination address in router solicitation messages is the all-routers multicast address with a scope of the link. When a router advertisement is sent in response to a router solicitation, the destination address in the router advertisement message is the unicast address of the source of the router solicitation message.

You can configure the following settings for router advertisement messages:

- The time interval between periodic router advertisement messages.
- The router lifetime value, which indicates the amount of time IPv6 nodes should consider the security appliance to be the default router.
- The IPv6 network prefixes used on the link.
- Whether or not an interface transmits router advertisement messages.

Unless otherwise noted, the router advertisement message settings are specific to an interface and are entered in interface configuration mode. See the following topics for information about changing these settings:

- Configuring the Router Advertisement Transmission Interval, page 14-6
- Configuring the Router Lifetime Value, page 14-6
- Suppressing Router Advertisement Messages, page 14-7

Configuring the Router Advertisement Transmission Interval

By default, router advertisements are sent out every 200 seconds. Valid values range from 3 to 1800 seconds.

The interval between transmissions should be less than or equal to the IPv6 router advertisement lifetime if the security appliance is configured as a default router. (See the "Configuring the Router Lifetime Value" section on page 14-6.) To prevent synchronization with other IPv6 nodes, randomly adjust the actual value used to within 20 percent of the desired value.

To change the interval between router advertisement transmissions on an interface, perform the following steps:

Step 1 Choose **Configuration** > **Device Setup** > **Interfaces**.

Step 2 Select the interface for which you want to configure the time.

The interface must have been configured with an IPv6 address. See the Chapter 9, "Configuring Interfaces," for more information.

Step 3 Click **Edit**.

The Edit Interface dialog box appears with three tabs: General, Advanced, and IPv6.

Step 4 Click the **IPv6 tab**.

Step 5 In the RA Interval field, enter a valid transmission interval value.

Note (Optional) To add a router advertisement transmission interval value in milliseconds instead of the default value in seconds, check the **RA Interval in Milliseconds** check box, and enter a value from 500 to 1800000 in the RA Interval field.

Step 6 Click **OK**.

Step 7 Click **Apply** to save the configuration.

Configuring the Router Lifetime Value

The router lifetime value specifies how long nodes on the local link should consider the security appliance as the default router on the link. Valid values range from 0 to 9000 seconds. The default is 1800 seconds. Entering 0 indicates that the security appliance should not be considered a default router on the selected interface.

To configure the router lifetime value in IPv6 router advertisements on an interface, perform the following steps:

Step 1 Choose **Configuration** > **Device Setup** > **Interfaces**.

Step 2 Select the interface for which you want to configure the lifetime value.

The interface must have been configured with an IPv6 address. See the Chapter 9, "Configuring Interfaces," for more information.

Step 3 Click **Edit**.

The Edit Interface dialog box appears with three tabs: General, Advanced, and IPv6.

Step 4 Click the **IPv6 tab**.

Step 5 In the RA Lifetime field, enter a valid lifetime value.

Step 6 Click **OK**.

Step 7 Click **Apply** to save the configuration.

Suppressing Router Advertisement Messages

By default, router advertisement messages are automatically sent in response to router solicitation messages. You may want to disable these messages on any interface for which you do not want the security appliance to supply the IPv6 prefix (for example, the outside interface).

To suppress IPv6 router advertisement transmissions on an interface, perform the following steps:

Step 1 Choose **Configuration** > **Device Setup** > **Interfaces**.

Step 2 Select the interface for which you want to configure the lifetime value. The interface must have been configured with an IPv6 address. See the "Configuring IPv6 Neighbor Discovery" section on page 14-2 for more information.

Step 3 Click **Edit**.

The Edit Interface dialog box appears with three tabs: General, Advanced, and IPv6.

Step 4 Click the **IPv6 tab**.

Step 5 Check the **Suppress RA** check box.

Step 6 Verify that the router advertisement message is suppressed on the interface that is configured for the IPv6 address.

Editing and Deleting Static Neighbors

This section includes the following topics:

- Editing Static Neighbors, page 14-8
- Deleting Static Neighbors, page 14-8

Editing Static Neighbors

To edit a static neighbor that is defined in your configuration, perform the following steps:

Step 1 Choose **Configuration** > **Device Management** > **Advanced** > **IPv6 Neighbor Discovery Cache**.

Step 2 In the main pane, select the neighbor you wish to edit, and click **Edit**.

The Edit IPv6 Static Neighbor dialog box appears.

Step 3 Enter all necessary changes, and click **OK**.

Note Before you apply the changes and permanently alter the neighbor in your configuration, you can click **Reset** to restore the original values.

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

Deleting Static Neighbors

To delete a static neighbor from your configuration, perform the following steps:

Step 1 Choose **Configuration** > **Device Management** > **Advanced** > **IPv6 Neighbor Discovery Cache**.

Step 2 In the main pane, select the neighbor you wish to delete, and click **Delete**.

The selected neighbor is removed from the list.

Note Before you apply the changes and permanently delete the neighbor from your configuration, you can click **Reset** to restore the original values.

Step 3 Click **Apply** to save the change to your current configuration.

Viewing and Clearing Dynamic Neighbors

When a host or node communicates with a neighbor, the neighbor is added to the neighbor discovery cache. The neighbor is removed from the cache when there is no longer any communication with the neighbor.

To view dynamically discovered neighbors and to clear neighbors from the IPv6 Neighbor Discovery Cache, perform the following steps:

Step 1 Choose **Monitoring** > **Interface Graphs** > **IPv6 Neighbor Discovery Cache**.

You can view all static and dynamically discovered neighbors from the IPv6 Neighbor Discovery Cache pane.

Step 2 To clear all dynamically discovered neighbors from the cache, click **Clear Dynamic Neighbor Entries**.

The neighbor information is removed from the cache.

Note This task clears only dynamically discovered neighbors from the cache. It does not clear static neighbors. To clear static neighbors, see the "Deleting Static Neighbors" section on page 14-8.

CHAPTER 15

DHCP, DNS and WCCP 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 hostnames) 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 hostnames at pre-defined intervals. DDNS allows frequently changing address-hostname 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 sections:

- DHCP Relay
- DHCP Server
- DNS Client
- Dynamic DNS
- WCCP

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 global 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 global 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 global server in the configuration or DHCP relay interface server.

Fields

- DHCP Relay Agent—*Display only.* Contains the fields for configuring the DHCP relay agent.
  - Interface—Displays the interface ID. Double-clicking the interface opens the Edit DHCP Relay Agent Settings dialog box, where you can enable the DHCP relay agent and configure the relay agent parameters.

    Note You can double-click anywhere in the row for a particular interface to open the dialog box for that interface.

  - 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 Global Server—Contains the fields for configuring the DHCP relay global servers.
  - 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 global 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 global server. You can define up to 4 DHCP relay global servers on the security appliance. This button is unavailable if you already have 4 DHCP relay global servers defined.
  - Edit—Opens the DHCP Relay - Edit DHCP Server dialog box, where you can edit the DHCP relay global server settings.
  - Delete—Removes the selected DHCP relay global server. The server is removed from the security appliance configuration when you apply or save your changes.
  - 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.

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

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 global 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 global 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 global 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
•	•	•	•	—

Add/Edit Global DHCP Relay Server

Define new global DHCP relay servers in the DHCP Relay - Add DHCP Server dialog box or edit existing server information in the DHCP Relay - Edit DHCP Server dialog box. You can define up to 4 DHCP relay global servers.

Restrictions

You cannot define a global 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

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 the 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.

  Note You can double-click anywhere in the row for a particular interface to open the dialog box for that 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.
- Global DHCP Options—Contains optional DHCP parameters.

  - Enable Auto-configuration from interface—Check this check box 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.
  - Allow VPN override—Overrides interface DHCP or PPPoE client WINS parameter with vpnclient parameter.
  - 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 this check box 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 this check box to specify that the DHCP server should update both the A and PTR RRs.
  - Override Client Settings—Check this check box 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

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 Auto-configuration on interface—Check this check box 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 this check box 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 this check box to specify that the DHCP server should update both the A and PTR RRs.
  - Override Client Settings—Check this check box 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

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

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 Clientless SSL VPN 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

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 Clientless SSL VPN configuration or certificate configuration. 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

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 hostnames at pre-defined intervals, DDNS allows frequently changing address-hostname 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

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

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
•	—	•	•	—

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

The Service Groups pane 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

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—Enter the dynamic service identifier, which means the service definition is dictated by the cache. Valid dynamic service numbers are 0 to 254, and are used as the name of the service group.
- Redirect List—Lets you choose the predefined access list that controls traffic redirected to this service group.
- Group List—Lets you choose the predefined access list that determines which web caches are allowed to participate in the service group. Only extended ACLs are allowed.
- Password—Enter a password up to seven characters long, which is used for MD5 authentication for messages received from the service group.
- Confirm Password—Reenter the password.
- Manage—Click to open the ACL Manager window, where you can create or change 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
•	•	•	•	•

Redirection

The Redirection pane 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

The Redirection pane lets you add or change packet redirection on the ingress of an interface using WCCP.

Fields

- Interface—Choose the interface on which to enable WCCP redirection.
- Service Group—Choose the service group.
- New—Opens the Add 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
•	•	•	•	•

CHAPTER 16

Configuring AAA Servers and the Local Database

This chapter describes support for AAA (pronounced “triple A”) and how to configure AAA servers and the local database.

This chapter includes the following sections:

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).

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.)

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.

This section includes the following topics:

About Authentication

Authentication controls access by requiring valid user credentials, which are typically a username and password. You can configure the security appliance to authenticate the following items:

- All administrative connections to the security appliance including the following sessions:
  - Telnet
  - SSH
  - Serial console
  - ASDM (using HTTPS)
  - VPN management access
- The **enable** command
- Network access
- VPN access

About Authorization

Authorization controls access *per user* after users authenticate. You can configure the security appliance to authorize the following items:

- Management commands
- Network access
- VPN access

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.

AAA Server and Local Database Support

The security appliance supports a variety of AAA server types and a local database that is stored on the security appliance. This section describes support for each AAA server type and the local database.

This section contains the following topics:

- Summary of Support, page 16-3
- RADIUS Server Support, page 16-4
- TACACS+ Server Support, page 16-4
- SDI Server Support, page 16-5
- NT Server Support, page 16-5
- Kerberos Server Support, page 16-5
- LDAP Server Support, page 16-6
- SSO Support for WebVPN with HTTP Forms, page 16-7
- Local Database Support, page 16-8

Summary of Support

Table 16-1 summarizes the support for each AAA service by each AAA server type, including the local database. For more information about support for a specific AAA server type, refer to the topics following the table.

Table 16-1 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
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.

RADIUS Server Support

The security appliance supports RADIUS servers.

This section contains the following topics:

- Authentication Methods, page 16-4
- Attribute Support, page 16-4
- RADIUS Authorization Functions, page 16-4

Authentication Methods

The security appliance supports the following authentication methods with RADIUS:

- PAP—For all connection types.
- CHAP—For L2TP-over-IPSec.
- MS-CHAPv1—For L2TP-over-IPSec.
- MS-CHAPv2—For L2TP-over-IPSec, and for regular IPSec remote access connections when the password management feature is enabled.

Attribute Support

The security appliance supports the following sets of RADIUS attributes:

- Authentication attributes defined in RFC 2138.
- Accounting attributes defined in RFC 2139.
- RADIUS attributes for tunneled protocol support, defined in RFC 2868.
- Cisco IOS VSAs, identified by RADIUS vendor ID 9.
- Cisco VPN-related VSAs, identified by RADIUS vendor ID 3076.
- Microsoft VSAs, defined in RFC 2548.

RADIUS Authorization Functions

The security appliance can use RADIUS servers for user authorization for network access using dynamic access lists or access list names per user. To implement dynamic access lists, you must configure the RADIUS server to support it. When the user authenticates, the RADIUS server sends a downloadable access list or access list name to the security appliance. Access to a given service is either permitted or denied by the access list. The security appliance deletes the access list when the authentication session expires.

TACACS+ Server Support

The security appliance supports TACACS+ authentication with ASCII, PAP, CHAP, and MS-CHAPv1.

SDI Server Support

The RSA SecurID servers are also known as SDI servers.

This section contains the following topics:

- SDI Version Support, page 16-5
- Two-step Authentication Process, page 16-5
- SDI Primary and Replica Servers, page 16-5

SDI Version Support

The security appliance supports SDI Version 5.0 and 6.0. SDI uses the concepts of an SDI primary and SDI replica servers. Each primary and its replicas share a single node secret file. The node secret file has its name based on the hexadecimal value of the ACE/Server IP address with .sdi appended.

A version 5.0 or 6.0 SDI server that you configure on the security appliance can be either the primary or any one of the replicas. See the "SDI Primary and Replica Servers" section on page 16-5 for information about how the SDI agent selects servers to authenticate users.

Two-step Authentication Process

SDI version 5.0 and 6.0 uses a two-step process to prevent an intruder from capturing information from an RSA SecurID authentication request and using it to authenticate to another server. The Agent first sends a lock request to the SecurID server before sending the user authentication request. The server locks the username, preventing another (replica) server from accepting it. This means that the same user cannot authenticate to two security appliances using the same authentication servers simultaneously. After a successful username lock, the security appliance sends the passcode.

SDI Primary and Replica Servers

The security appliance obtains the server list when the first user authenticates to the configured server, which can be either a primary or a replica. The security appliance then assigns priorities to each of the servers on the list, and subsequent server selection derives at random from those assigned priorities. The highest priority servers have a higher likelihood of being selected.

NT Server Support

The security appliance supports Microsoft Windows server operating systems that support NTLM version 1, collectively referred to as NT servers.

Note NT servers have a maximum length of 14 characters for user passwords. Longer passwords are truncated. This is a limitation of NTLM version 1.

Kerberos Server Support

The security appliance supports 3DES, DES, and RC4 encryption types.

Note The security appliance does not support changing user passwords during tunnel negotiation. To avoid this situation happening inadvertently, disable password expiration on the Kerberos/Active Directory server for users connecting to the security appliance.

LDAP Server Support

This section describes LDAP server support, and includes the following topics:

- Authentication with LDAP, page 16-6
- Securing LDAP Authentication with SASL, page 16-6
- LDAP Server Types, page 16-7
- Authorization with LDAP for VPN, page 16-7

Authentication with LDAP

During authentication, the security appliance acts as a client proxy to the LDAP server for the user, and authenticates to the LDAP server in either plain text or using the Simple Authentication and Security Layer (SASL) protocol. By default, the security appliance passes authentication parameters, usually a username and password, to the LDAP server in plain text. Whether using SASL or plain text, you can secure the communications between the security appliance and the LDAP server with SSL.

Note If you do not configure SASL, we strongly recommend that you secure LDAP communications with SSL.

When user LDAP authentication has succeeded, the LDAP server returns the attributes for the authenticated user. For VPN authentication, these attributes generally include authorization data which is applied to the VPN session. Thus, using LDAP accomplishes authentication and authorization in a single step.

Securing LDAP Authentication with SASL

The security appliance supports the following SASL mechanisms, listed in order of increasing strength:

- Digest-MD5 — The security appliance responds to the LDAP server with an MD5 value computed from the username and password.
- Kerberos — The security appliance responds to the LDAP server by sending the username and realm using the GSSAPI (Generic Security Services Application Programming Interface) Kerberos mechanism.

You can configure the security appliance and LDAP server to support any combination of these SASL mechanisms. If you configure multiple mechanisms, the security appliance retrieves the list of SASL mechanisms configured on the server and sets the authentication mechanism to the strongest mechanism configured on both the security appliance and the server. For example, if both the LDAP server and the security appliance support both mechanisms, the security appliance selects Kerberos, the stronger of the mechanisms.

LDAP Server Types

The security appliance supports LDAP version 3 and is compatible with the Sun Microsystems JAVA System Directory Server (formerly named the Sun ONE Directory Server), the Microsoft Active Directory, and other LDAPv3 directory servers.

By default, the security appliance auto-detects whether it is connected to a Microsoft Active Directory, a Sun LDAP directory server, or a generic LDAPv3 directory server. However, if auto-detection fails to determine the LDAP server type, and you know the server is either a Microsoft, Sun or generic LDAP server, you can manually configure the server type.

Note

- Sun—The DN configured on the security appliance to access a Sun directory server must be able to access the default password policy on that server. We recommend using the directory administrator, or a user with directory administrator privileges, as the DN. Alternatively, you can place an ACI on the default password policy.
- Microsoft—You must configure LDAP over SSL to enable password management with Microsoft Active Directory.
- Generic—The security appliance does not support password management with a generic LDAPv3 directory server.

Authorization with LDAP for VPN

When user LDAP authentication for VPN access has succeeded, the security appliance queries the LDAP server which returns LDAP attributes. These attributes generally include authorization data that applies to the VPN session. Thus, using LDAP accomplishes authentication and authorization in a single step.

There may be cases, however, where you require authorization from an LDAP directory server that is separate and distinct from the authentication mechanism. For example, if you use an SDI or certificate server for authentication, no authorization information is passed back. For user authorizations in this case, you can query an LDAP directory after successful authentication, accomplishing authentication and authorization in two steps.

SSO Support for WebVPN with HTTP Forms

The security appliance can use the HTTP Form protocol for single sign-on (SSO) authentication of WebVPN users only. Single sign-on support lets WebVPN users enter a username and password only once to access multiple protected services and Web servers. 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. The security appliance keeps this cookie on behalf of the user and uses it to authenticate the user to secure websites within the domain protected by the SSO server.

In addition to the HTTP Form protocol, WebVPN administrators can choose to configure SSO with the HTTP Basic and NTLM authentication protocols (the **auto-signon** command), or with Computer Associates eTrust SiteMinder SSO server (formerly Netegrity SiteMinder) as well. For an in-depth discussion of configuring SSO with either HTTP Forms, **auto-signon** or SiteMinder, see the Clientless SSL VPN chapter.

Local Database Support

The security appliance maintains a local database that you can populate with user profiles.

This section contains the following topics:

- User Profiles, page 16-8
- Fallback Support, page 16-8

User Profiles

User profiles contain, at a minimum, a username. Typically, a password is assigned to each username, although passwords are optional.

The **username attributes** command lets you enter the username mode. In this mode, you can add other information to a specific user profile. The information you can add includes VPN-related attributes, such as a VPN session timeout value.

Fallback Support

The local database can act as a fallback method for several functions. 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.

The local database supports the following fallback functions:

- **Console and enable password authentication**—When you use the **aaa authentication console** command, you can add the **LOCAL** keyword after the AAA server group tag. If the servers in the group all are unavailable, the security appliance uses the local database to authenticate administrative access. This can include enable password authentication, too.
- **Command authorization**—When you use the **aaa authorization command** command, you can add the **LOCAL** keyword after the AAA server group tag. If the TACACS+ servers in the group all are unavailable, the local database is used to authorize commands based on privilege levels.
- **VPN authentication and authorization**—VPN authentication and authorization are supported to enable remote access to the security appliance if AAA servers that normally support these VPN services are unavailable. The **authentication-server-group** command, available in tunnel-group general attributes mode, lets you specify the **LOCAL** keyword when you are configuring attributes of a tunnel group. When VPN client of an administrator specifies a tunnel group configured to fallback to the local database, the VPN tunnel can be established even if the AAA server group is unavailable, provided that the local database is configured with the necessary attributes.

Configuring AAA Server Groups

If you want to use an external AAA server for authentication, authorization, or accounting, you must first create at least one AAA server group per AAA protocol and add one or more servers to each group. You identify AAA server groups by name. Each server group is specific to one type of server: Kerberos, LDAP, NT, RADIUS, SDI, or TACACS+.

You can have up to 100 server groups in single mode or 4 server groups per context in multiple mode. Each group can have up to 16 servers in single mode or 4 servers in multiple mode. When a user logs in, the servers are accessed one at a time starting with the first server you specify in the configuration, until a server responds. If all servers in the group are unavailable, the security appliance tries the local database if you configured it as a fallback method (management authentication and authorization only). If you do not have a fallback method, the security appliance continues to try the AAA servers.

This section includes the following procedures:

- Adding a Server Group, page 16-9
- Adding a Server to a Group, page 16-10
- AAA Server Parameters, page 16-11

Adding a Server Group

To add a server group, perform the following steps:

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Server Groups**.

Step 2 In the AAA Server Groups area, click **Add**.

The Add AAA Server Group dialog box appears.

Step 3 In the Server Group field, add a name for the group.

Step 4 From the Protocol drop-down list, choose the server type:

- **RADIUS**
- **TACACS+**
- **SDI**
- **NT Domain**
- **Kerberos**
- **LDAP**
- **HTTP Form**

Step 5 In the Accounting Mode field, click the radio button for the mode you want to use (**Simultaneous** or **Single**).

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.

Note This option is not available for the following protocols: HTTP Form, sdi, NT, Kerberos, and LDAP.

Step 6 In the Reactivation Mode field, click the radio button for the mode you want to use (**Depletion** or **Timed**).

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.

Step 7 If you chose Depletion reactivation mode, add a time interval in the Dead Time field.

The Dead Time is the duration of time, in minutes, to elapse between the disabling of the last server in a group and the subsequent reenabling of all servers.

Step 8 In the Max Failed Attempts field, add the number of failed attempts permitted.

This option sets the number of failed connection attempts allowed before declaring a nonresponsive server to be inactive.

Step 9 (Optional) If you are adding a RADIUS server type, you can perform the following optional steps:

a. Check the **Enable interim accounting update** check box if you want to enable multi-session accounting for SSL clientless and AnyConnect sessions.

b. Click **VPN3K Compatibility Option** to expand the list, and click one of the following radio buttons to specify whether or not a downloadable ACL received from RADIUS should be merged with a Cisco AV-pair ACL:

   – **Do not merge**

   – **Place the downloadable ACL after Cisco AV-pair ACL**

   – **Place the downloadable ACL before Cisco AV-pair ACL**

Step 10 Click **OK**.

The dialog box closes, and the server group is added to the AAA server groups table.

Step 11 In the AAA Server Groups dialog box, click **Apply** to save the changes.

The changes are saved.

Adding a Server to a Group

To add a AAA server to a group, perform the following steps:

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Server Groups**, and in the AAA Server Groups area, click the server group to which you want to add a server.

The row is highlighted in the table.

Step 2 In the Servers in the Selected Group area (lower pane), click **Add**.

The Add AAA Server Group dialog box appears for the server group.

Step 3 From the Interface Name drop-down menu, choose the interface name where the authentication server resides.

Step 4 In the Server Name or IP Address field, add either a server name or IP address for the server you are adding to the group.

Step 5 In the Timeout field, either add a timeout value or keep the default. The timeout is the duration of time, in seconds, that the security appliance waits for a response from the primary server before sending the request to the backup server.

Step 6 The other parameters available depend on the server type. See the following sections for parameters unique to each server type:

- RADIUS Server Fields, page 16-11
- TACACS+ Server Fields, page 16-13
- SDI Server Fields, page 16-13
- Windows NT Domain Server Fields, page 16-13
- Kerberos Server Fields, page 16-14
- LDAP Server Fields, page 16-15
- HTTP Form Server Fields, page 16-17

Step 7 Click **OK**.

The dialog box closes and the AAA server is added to the AAA server group.

Step 8 In the AAA Server Groups pane, click **Apply** to save the changes.

The changes are saved.

AAA Server Parameters

The following sections list the unique fields for each server type when adding a server to a server group (see the "Adding a Server to a Group" section on page 16-10):

- RADIUS Server Fields, page 16-11
- TACACS+ Server Fields, page 16-13
- SDI Server Fields, page 16-13
- Windows NT Domain Server Fields, page 16-13
- Kerberos Server Fields, page 16-14
- LDAP Server Fields, page 16-15
- HTTP Form Server Fields, page 16-17

RADIUS Server Fields

The following table describes the unique fields for configuring RADIUS servers, for use with the "Adding a Server to a Group" section on page 16-10.

Field	Description
Server Authentication Port	The server port to be used for authentication of users. The default port is 1645.
Server Accounting Port	The server port to be used for accounting of users. The default port is 1646.
Retry Interval	The duration of time, 1 to 10 seconds, that the security appliance waits between attempts to contact the server.

Field	Description
Server Secret Key	The shared secret key used to authenticate the RADIUS server to the security appliance. The server secret you configure here should match the one configured on the RADIUS server. If you do not know the server secret, ask the RADIUS server administrator. The maximum field length is 64 characters.
Common Password	A case-sensitive password that is common among users who access this RADIUS **authorization** server through this security appliance. Be sure to provide this information to your RADIUS server administrator. **Note** For an authentication RADIUS server (rather than authorization) do not configure a common password. If you leave this field blank, the users username is the password for accessing this RADIUS **authorization** server. Never use a RADIUS authorization server for authentication. Common passwords or usernames as passwords are less secure than assigning unique user passwords. **Note** Although the password is required by the RADIUS protocol and the RADIUS server, users do not need to know it.
ACL Netmask Convert	How you want the security appliance to handle netmasks received in downloadable access lists. • 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; **Note** Because some wildcard expressions are difficult to detect clearly, this setting may misinterpret a wildcard netmask expression as a standard netmask expression. • 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. • 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.
Microsoft CHAPv2 Capable	If you use double authentication and enable password management in the tunnel group, then the primary and secondary authentication requests include MS-CHAPv2 request attributes. If a RADIUS server does not support MS-CHAPv2, then you can configure that server to send a non-MS-CHAPv2 authentication request by unchecking this check box.

TACACS+ Server Fields

The following table describes the unique fields for configuring TACACS+ servers, for use with the "Adding a Server to a Group" section on page 16-10.

Field	Description
Server Port	The port to be used for this server.
Server Secret Key	The shared secret key used to authenticate the TACACS+ server to the security appliance. The server secret you configure here should match the one configured on the TACACS+ server. If you do not know the server secret, ask the RADIUS server administrator. The maximum field length is 64 characters.

SDI Server Fields

The following table describes the unique fields for configuring SDI servers, for use with the "Adding a Server to a Group" section on page 16-10.

Field	Description
Server Port	The TCP port number by which this server is accessed.
Retry Interval	The duration of time, 1 to 10 seconds, that the security appliance waits between attempts to contact the server.

Windows NT Domain Server Fields

The following table describes the unique fields for configuring Windows NT Domain servers, for use with the "Adding a Server to a Group" section on page 16-10.

Field	Description
Server Port	Port number 139, or the TCP port number used by the security appliance to communicate with the Windows NT server.
Domain Controller	The host name (no more than 15 characters) of the NT Primary Domain Controller for this server. For example, PDC01. You must enter a name, and it must be the correct host name for the server whose IP Address you added in the field, Authentication Server Address. If the name is incorrect, authentication fails.

Kerberos Server Fields

The following table describes the unique fields for configuring Kerberos servers, for use with the "Adding a Server to a Group" section on page 16-10.

Field	Description
Server Port	Server port number 88, or the UDP port number over which the security appliance communicates with the Kerberos server.
Retry Interval	The duration of time, 1 to 10 seconds, that the security appliance waits between attempts to contact the server.
Realm	The name of the Kerberos realm. For example: • EXAMPLE.COM • EXAMPLE.NET • EXAMPLE.ORG **Note** Most Kerberos servers require the realm to be all uppercase for authentication to succeed. 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 • 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 Server Fields

The following table describes the unique fields for configuring LDAP servers, for use with the "Adding a Server to a Group" section on page 16-10.

Field	Description
Enable LDAP over SSL check box	When checked, SSL secures communications between the security appliance and the LDAP server. Also called secure LDAP (LDAP-S). **Note** If you do not configure SASL protocol, we strongly recommend that you secure LDAP communications with SSL.
Server Port	TCP port number 389, the port which the security appliance uses to access the LDAP server for simple (non-secure) authentication or TCP port 636 for secure authentication (LDAP-S). All LDAP servers support authentication and authorization. Only Microsoft AD and Sun LDAP servers additionally provide VPN remote access password management capability, which requires LDAP-S.
Server type	A drop-down list for choosing one of the following LDAP server types: • Detect Automatically/Use Generic Type • Microsoft • Novell • OpenLDAP • Sun
Base DN	The Base Distinguished Name (DN), or location in the LDAP hierarchy where the server should begin searching when it receives an LDAP request. For example, OU=people, dc=cisco, dc=com.
Scope	The extent of the search the server should make in the LDAP hierarchy when it receives an authorization request. The available options are: • One Level: Searches only one level beneath the Base DN. This option is quicker. • All Levels: Searches all levels beneath the Base DN; in other words, search the entire subtree hierarchy. This option takes more time.
Naming Attribute(s)	The Relative Distinguished Name attribute (or attributes) that uniquely identifies an entry on the LDAP server. Common naming attributes are Common Name (CN), sAMAccountName, userPrincipalName, and User ID (uid).

Field	Description
Login DN	The security appliance uses the Login Distinguished Name (DN) and Login Password to establish trust (bind) with an LDAP server. The Login DN represents a user record in the LDAP server that the administrator uses for binding. When binding, the security appliance authenticates to the server using the Login DN and the Login Password. When performing a Microsoft Active Directory read-only operation (such as for authentication, authorization, or group-search), the security appliance can bind with a Login DN with less privileges. For example, the Login DN can be a user whose AD “Member Of” designation is part of Domain Users. For VPN password management operations, the Login DN needs elevated privileges and must be part of the Account Operators AD group. An example of a Login DN include: `cn=Binduser1,ou=Admins,ou=Users,dc=company_A,dc=com` The security appliance supports: • Simple LDAP authentication with an unencrypted password on port 389 • Secure LDAP (LDAP-S) on port 636 • Simple Authentication and Security Layer (SASL) MD5 • SASL Kerberos. The security appliance does not support anonymous authentication.
Login Password	The password for the Login DN user account. The characters you type are replaced with asterisks.
LDAP Attribute Map	The LDAP attribute maps that you can apply to LDAP server. Used to map Cisco attribute names to user-defined attribute names and values. See the “Configuring LDAP Attribute Maps” section on page 16-22.
SASL MD5 authentication check box	When checked, the MD5 mechanism of the Simple Authentication and Security Layer (SASL) authenticates communications between the security appliance and the LDAP server.
SASL Kerberos authentication	When checked, the Kerberos mechanism of the SASL secures authentication communications between the security appliance and the LDAP server.
Kerberos Server Group	The Kerberos server or server group used for authentication. The Kerberos Server group option is disabled by default and is enabled only when SASL Kerberos authentication is chosen.

Field	Description
Group Base DN	Used only for Active Directory servers using LDAP protocol. This DN specifies the location in the LDAP hierarchy to begin searching for the AD groups. That is, the list of memberOf enumerations. If this field is not configured, the security appliance uses the Base DN for AD group retrieval. ASDM uses the list of retrieved AD groups to define AAA selection criterion for dynamic access policies. For more information, see the **show ad-groups** command in CLI Command Reference Guide.
Group Search Timeout	Specifies the maximum time to wait for a response from an Active Directory server queried for available groups.

HTTP Form Server Fields

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.

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 appliancee. See the *Cisco ASA 5500 Series Configuration Guide using the CLI*, for more detail on extracting these parameters from the HTTP exchanges.

The following table describes the unique fields for configuring HTTP Form servers, for use with the "Adding a Server to a Group" section on page 16-10.

Field	Description
Start URL	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	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	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	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.

Field	Description
Hidden Values	The 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) 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.

Testing Server Authentication and Authorization

To determine whether the security appliance can contact an AAA server and authenticate or authorize a user, perform the following steps:

Step 1 From the Configuration > Device Management > Users/AAA > AAA Server Groups > AAA Server Groups table, click the server group where the server resides.

The row is highlighted in the table.

Step 2 From the Servers in the Selected Group table, click the server you want to test.

The row is highlighted in the table.

Step 3 Click **Test**.

The Test AAA Server dialog box appears for that server.

Step 4 Click the type of test you want to perform, **Authentication** or *Authorization*.

Step 5 In the Username field, add a username.

Step 6 If you are testing authentication, in the Password field, add the password for the username.

Step 7 Click **OK**.

The security appliance sends an authentication or authorization test message to the server. If the test fails, ASDM displays an error message.

Adding a User Account

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 Configuring the Device Name and Password, page 8-15). 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, 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
- Telnet and SSH authentication
- enable command authentication

  This setting is for CLI-access only and does not affect the ASDM login.
- Command authorization

  If you turn on 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.
- 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 rules that use the local database in the system execution space.

To add a user account to the security appliance local database, perform the following steps:

Step 1 From the Configuration > Device Management > Users/AAA > User Accounts pane, click **Add**.

The Add User Account—Identity dialog box appears.

Step 2 In the Username field, add a username between 4 to 64 characters long.

Step 3 In the Password field add a password between 3 and 32 characters. Entries are case-sensitive. The field displays only asterisks. To protect security, we recommend a password length of at least 8 characters.

Step 4 In the Confirm Password field, add the password again.

For security purposes, only asterisks appear in the password fields.

Step 5 To enable MSCHAP authentication, check **User authenticated using MSCHAP**.

This option specifies that the password is converted to unicode and hashed using MD4 after you enter it. Use this feature if users are authenticated using MSCHAPv1 or MSCHAPv2.

Step 6 To specify the VPN groups that the user belongs to, enter a group name in the Member of field, and click **Add**.

To delete a VPN group, choose the group in the window, and click **Delete**.

Step 7 In the Access Restriction area, set the management access level for a user. You must first enable management authorization using the **Perform authorization for exec shell access** option on the Configuration > Device Management > Users/AAA > AAA Access > Authorization tab.

Choose one of the following options:

- **Full Access (ASDM, Telnet, SSH and console)**—If you configure authentication for management access using the local database , then this option lets the user use ASDM, SSH, Telnet, and the console port. If you also configure enable authentication, then the user can access global configuration mode.
  - **Privilege Level**—Selects the privilege level for this user to use with local command authorization. The range is 0 (lowest) to 15 (highest) .
- **CLI login prompt for SSH, Telnet and console (no ASDM access)**—If you configure authentication for management access using the local database , then this option lets the user use SSH, Telnet, and the console port. The user cannot use ASDM for configuration (if you configure HTTP authentication). ASDM monitoring is allowed. If you also configure enable authentication, then the user cannot access global configuration mode.
- **No ASDM, SSH, Telnet, or console access**—If you configure authentication for management access using the local database, then this option disallows the user from accessing any management access method for which you configured authentication (excluding the Serial option; serial access is allowed).

Step 8 If you want to configure VPN policy attributes for this user, see the "Configuring VPN Policy Attributes for a User" section on page 16-20.

Step 9 Click **Apply**.

The user is added to the local security appliance database and changes are saved to the running configuration.

Note To configure the enable password from the User Accounts pane (see Configuring the Device Name and Password, page 8-15), change the password for the enable_15 user. The enable_15 user is always present in this pane, 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.

Configuring VPN Policy Attributes for a User

By default, each user inherits the settings set in the VPN policy. To override the settings, you can customize VPN attributes by performing the following steps:

Step 1 If you have not already added a user according to the "Adding a User Account" section on page 16-18, from the Configuration > Device Management > Users/AAA > User Accounts pane, click **Add**.

The Add User Account—Identity dialog box appears.

Step 2 In the left-hand pane, click **VPN Policy**.

By default, the Inherit check box is checked for each option, which means the user account inherits the settings from the VPN policy. To override each setting, uncheck **Inherit**, and fill in a new value:

- Group Policy—Choose a group policy from the list.
- 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.
  - Clientless SSL VPN—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. Clientless SSL VPN 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.
  - SSL VPN Client—Lets users connect after downloading the Cisco AnyConnect Client application. Users use a clientless SSL VPN connection to download this application the first time. Client updates then occur automatically as needed whenever the user connects.
  - 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.

- IPv4 Filter/IPv6 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 pane.
- Manage—Displays the ACL Manager pane, 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.

Step 3 To change Connection Settings, uncheck **Inherit**, and fill in a new value:

- 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 users of clientless SSL VPN connections.

Step 4 To set a dedicated IP address for this user, enter an IP address and subnet mask in the Dedicated IP Address (Optional) area.

Step 5 To configure clientless SSL settings, in the left-hand pane, click **Clientless SSL VPN**.

To override each setting, uncheck **Inherit**, and fill in a new value. See the "Group Policies" section on page 36-5.

Step 6 To configure SSL VPN settings, in the left-hand pane, click **SSL VPN Client**.

To override each setting, uncheck **Inherit**, and fill in a new value. See the "Configuring AnyConnect (SSL) VPN Client Connections" section on page 36-41.

Step 7 Click **Apply**.

Configuring LDAP Attribute Maps

If you are introducing a security appliance to an existing LDAP directory, your existing LDAP attribute names and values are probably different from the existing ones. You must create LDAP attribute maps that map your existing user-defined attribute names and values to Cisco attribute names and values that are compatible with the security appliance. You can then bind these attribute maps to LDAP servers or remove them as needed. You can also show or clear attribute maps.

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.

The names of frequently mapped Cisco LDAP attributes and the type of user-defined attributes they would commonly be mapped to include:

```
IETF-Radius-Class — Department or user group
IETF-Radius-Filter-Id — Access control list
IETF-Radius-Framed-IP-Address — A static IP address
IPSec-Banner1 — A organization title
Tunneling-Protocols — Allow or deny dial-in
```

For a list of Cisco LDAP attribute names and values, see Appendix C, "Configuring an External LDAP Server".

To map the LDAP attribute names used in your organization to their Cisco counterparts on the security appliance, perform the following steps:

Step 1 From the Configuration > Remote Access VPN > AAA Local Users > LDAP Attribute Map pane, click **Add**.

The Add LDAP Attribute Map dialog box appears with the Map Name tab active.

Step 2 In the Name field, add a name for the map.

Step 3 In the Customer Name field, add the name of your organization's corresponding attribute.

Step 4 From the Cisco Name drop-down list, choose an attribute.

Step 5 Click **Add**.

Step 6 To add more names, repeat steps 1 through 5.

Step 7 To map the customer names, click the **Map Value** tab.

Step 8 Click **Add**.

The Add LDAP Attributes Map Value dialog box appears.

Step 9 Choose the attribute from the Customer Name drop-down list.

Step 10 In the Customer Value field, add the value for this attribute.

Step 11 In the Cisco Value field, add the Cisco value that the value in step 10 maps to.

Step 12 Click **Add**.

The values are mapped.

Step 13 To map more names, repeat steps 8 through 12.

Step 14 Click **OK** to return to the Map Value tab, and then click **OK** again to close the dialog box.

Step 15 In the LDAP Attribute Map pane, click **Apply**.

The value mappings are saved in the running configuration.

Adding an Authentication Prompt

You can 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 you do not specify an authentication prompt, users will see the following when authenticating with a RADIUS or TACACS+ server:

Connection type	Default prompt
FTP	`FTP authentication`
HTTP	`HTTP Authentication`
Telnet	None

To add an authentication prompt, perform the following steps:

Step 1 From the Configuration > Device Management > Users/AAA > Authentication Prompt pane, add a message to appear above the username and password prompts that users see when logging in by entering text in the Prompt field.

The following are maximum characters allowed for authentication prompts:

Application	Character limit for Authentication prompt
Microsoft Internet Explorer	37
Telnet	235
FTP	235

Step 2 In the Messages area, add messages in the User accepted message and User rejected message fields.

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.

Step 3 Click **Apply**.

The changes are saved to the running configuration.

CHAPTER 17

High Availability

This chapter provides an overview of failover and shows how to configure failover with high availability.

This chapter contains the following sections:

- Understanding Failover, page 17-1
- Configuring Failover with the High Availability and Scalability Wizard, page 17-4
- Field Information for the Failover Panes, page 17-15

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. 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.

Note For failover to function properly, we recommend that you configure at least one management interface with the standby IP address.

For more information about the types of failover, see the following topics:

- Active/Standby Failover, page 17-2
- Active/Active Failover, page 17-2
- Stateless (Regular) Failover, page 17-3
- Stateful Failover, page 17-4

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.

Note During a successful failover event on the security appliance, the interfaces are brought down, roles are switched (IP addresses and MAC addresses are swapped), and the interfaces are brought up again. However, the process is transparent to users. The security appliance does not send link-down messages or system log messages to notify users that interfaces were taken down during failover (or link-up messages for interfaces brought up by the failover process).

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.

Note During a successful failover event on the security appliance, the interfaces are brought down, roles are switched (IP addresses and MAC addresses are swapped), and the interfaces are brought up again. However, the process is transparent to users. The security appliance does not send link-down messages or system log messages to notify users that interfaces were taken down during failover (or link-up messages for interfaces brought up by the failover process).

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. 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 IPv6 Neighbor Discovery 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 walks you through a step-by-step 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, page 17-5
- Configuring Active/Active Failover with the High Availability and Scalability Wizard, page 17-5
- Configuring Active/Standby Failover with the High Availability and Scalability Wizard, page 17-6
- Configuring VPN Load Balancing with the High Availability and Scalability Wizard, page 17-7

- Field Information for the High Availability and Scalability Wizard, page 17-7

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 **Next**. 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 **Back**. 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.

Note During a successful failover event on the security appliance, the interfaces are brought down, roles are switched (IP addresses and MAC addresses are swapped), and the interfaces are brought up again. However, the process is transparent to users. The security appliance does not send link-down messages or system log messages to notify users that interfaces were taken down during failover (or link-up messages for interfaces brought up by the failover process).

Each step in the following 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 in the Choose the type of failover configuration pane.

See Choose the Type of Failover Configuration, page 17-8 for more information about this screen.

Step 2 Enter the IP address of the failover peer in 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, page 17-9 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 in 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, page 17-9 for more information about this screen.

Step 4 Assign security contexts to failover groups in the Context Configuration screen. You can add and delete contexts on this screen.

See Security Context Configuration, page 17-10 for more information about this screen.

Step 5 Define the Failover Link in the Failover Link Configuration screen.

See Failover Link Configuration, page 17-10 for more information about this screen.

Step 6 (Not available on the ASA 5505 security appliance) Define the Stateful Failover link in the State Link Configuration screen.

See State Link Configuration, page 17-11 for more information about this screen.

Step 7 Add standby addresses to the security appliance interfaces in the Standby Address Configuration screen.

See Standby Address Configuration, page 17-12 for more information about this screen.

Step 8 Review your configuration in the Summary screen. If necessary, use the Back button to go to a previous screen and make changes.

See Summary, page 17-14 for more information about this screen.

Step 9 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.

Note During a successful failover event on the security appliance, the interfaces are brought down, roles are switched (IP addresses and MAC addresses are swapped), and the interfaces are brought up again. However, the process is transparent to users. The security appliance does not send link-down messages or system log messages to notify users that interfaces were taken down during failover (or link-up messages for interfaces brought up by the failover process).

Each step in the following 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, page 17-8 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, page 17-9 for more information about this screen.

Step 3 Define the Failover Link in the Failover Link Configuration screen.

See Failover Link Configuration, page 17-10 for more information about this screen.

Step 4 (Not available on the ASA 5505 security appliance) Define the Stateful Failover link in the State Link Configuration screen.

See State Link Configuration, page 17-11 for more information about this screen.

Step 5 Add standby addresses to the security appliance interfaces in the Standby Address Configuration screen.

See Standby Address Configuration, page 17-12 for more information about this screen.

Step 6 Review your configuration in the Summary screen. If necessary, use the Back button to go to a previous screen and make changes.

See Summary, page 17-14 for more information about this screen.

Step 7 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 in the Choose the type of failover configuration screen.

See Choose the Type of Failover Configuration, page 17-8 for more information about this screen.

Step 2 Configure the VPN load balancing settings in the VPN Cluster Load Balancing Configuration screen.

See VPN Cluster Load Balancing Configuration, page 17-13 for more information about this screen.

Step 3 Review your configuration in the Summary screen. If necessary, use the Back button to go to a previous screen and make changes.

See Summary, page 17-14 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.

This section contains the following topics:

- Choose the Type of Failover Configuration, page 17-8
- Check Failover Peer Connectivity and Compatibility, page 17-9
- Change Device to Multiple Mode, page 17-9

- Security Context Configuration, page 17-10
- Failover Link Configuration, page 17-10
- State Link Configuration, page 17-11
- Standby Address Configuration, page 17-12
- VPN Cluster Load Balancing Configuration, page 17-13
- Summary, page 17-14

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
•	•	•	—	•

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 choose 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 allows you to define the failover link.

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 choose a subnet mask for the Active IP and Standby IP addresses.
- Switch Port—(ASA 5505 only) Choose the switch port from the drop-down list, which includes the current VLAN assigned to each switch port and any name associated with the VLAN. Because there is a default VLAN for every switch port, we recommend that you not choose VLAN 1 for inside, as you will have one less inside port for another use.

Note To provide sufficient brandwidth for failover, we discurage using trunks or PoE for failover.

- 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.
- 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—Choose 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 To use VPN load balancing, you must have an ASA Model 5510 with a Plus license or an ASA Model 5520 or 5540. VPN load balancing also 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 the following:

- 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 (Version 3.0 and later), the Cisco VPN 3002 hardware client (Version 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

- 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 check 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, check 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 check 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.

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.

- 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.
- Send FQDN to client—Check this check box to cause the VPN cluster master to send a fully qualified domain name using the host and domain name of the cluster device instead of the outside IP address when redirecting VPN client connections to that cluster device.

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 appears in 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 includes the following topics:

- Failover - Single Mode
- Failover-Multiple Mode, Security Context
- Failover-Multiple Mode, System

Failover - Single Mode

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

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 changes.

  – Peer IP Address—Enter an IP address on the peer unit that ASDM can connect to. This field appears in the Do you want to configure the failover peer firewall dialog box. You can enter IPv4 or IPv6 IP addresses.

- 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.

- 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. The IP address can be an IPv4 or an IPv6 address.

  – Subnet Mask/Prefix Length—Depending upon the type of address specified for the Active IP, enter a subnet mask (IPv4 addresses) or a prefix length (IPv6 address) 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. The IP address can be an IPv4 or an IPv6 address.

  – 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/Prefix Length, 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 chosen from the Interface drop-down list.
- Subnet Mask/Prefix Length—Specifies the mask (IPv4 address) or prefix (IPv6 address) 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 chose 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 chosen from 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 chosen from the Interface drop-down list.
- 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: Interfaces (Routed Firewall Mode)

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)

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/Prefix Length—Identifies the mask (IPv4 address) or prefix (IPv6 address) 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)

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)

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

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

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

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 includes the following topics:

- Failover - Routed
- Failover - Transparent

Failover - Routed

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

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/Prefix Length—Identifies the mask (for IPv4 addresses) or prefix (for IPv6 addresses) 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

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

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

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

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.

Note During a successful failover event on the security appliance, the interfaces are brought down, roles are switched (IP addresses and MAC addresses are swapped), and the interfaces are brought up again. However, the process is transparent to users. The security appliance does not send link-down messages or system log messages to notify users that interfaces were taken down during failover (or link-up messages for interfaces brought up by the failover process).

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.
- 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. The IP address can be an IPv4 or an IPv6 address.
  - Subnet Mask/Prefix Length—Depending upon the type of address specified for the Active IP, enter a subnet mask (IPv4 addresses) or a prefix length (IPv6 address) 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. The IP address can be an IPv4 or an IPv6 address.
  - 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 subinterfaces 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 primary unit. This field is dimmed if the LAN Failover interface or Use Named option is chosen from the Interface drop-down list.
- Subnet Mask/Prefix Length—Specifies the mask (IPv4 address) or prefix (IPv6 address) 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 chose 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 chosen from 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 chosen from the Interface drop-down list.
- 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

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

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 During a successful failover event on the security appliance, the interfaces are brought down, roles are switched (IP addresses and MAC addresses are swapped), and the interfaces are brought up again. However, the process is transparent to users. The security appliance does not send link-down messages or system log messages to notify users that interfaces were taken down during failover (or link-up messages for interfaces brought up by the failover process).

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

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

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

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

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 18

Configuring Management Access

This chapter contains the following topics:

Note To configure the Management IP address for transparent firewall mode, see the "Configuring the Management IP Address for Transparent Firewall Mode" section on page 8-1.

Configuring Device Access for ASDM, Telnet, or SSH

This section describes how to allow clients to access the device using ASDM, Telnet, or SSH. To configure access to the security appliance, perform the following steps:

Step 1 From the Configuration > Device Management > Management Access > ASDM/HTTPS/Telnet/SSH pane, click **Add**.

The Add Device Access Configuration dialog box appears in the right-hand pane.

Step 2 Choose the type of session from the three options listed: ASDM/HTTPS, Telnet, or SSH.

Step 3 From the Interface Name drop-down list, choose the interface to use for administrative access.

Step 4 In the IP Address field, add the IP address of the network or host that is allowed access. The field allows IPv6 addresses.

Note When you enter a colon (:) in the IP Address field for an IPv6 address, the Netmask field changes to Prefix Length.

Step 5 From the Mask drop-down list, choose the mask associated with the network or host that is allowed access.

Step 6 For ASDM/HTTPS sessions, verify that the Enable HTTP Server check box is checked. This is the default setting.

Step 7 Specify the port number. The default port is 443.

Step 8 Adjust the Idle Timeout or Session Timeout if necessary. There is no timeout value by default. This setting is available only in single, routed mode.

Step 9 For Telnet sessions, the default timeout value is 5 minutes. To change this value, type a new one in the Telnet Timeout field.

Step 10 For SSH sessions, the default timeout value is 5 minutes. To change this value, type a new one in the SSH Timeout field.

Step 11 Click **Apply**.

The changes are saved to the running configuration.

Configuring CLI Parameters

This section includes the following topics:

- Adding a Banner, page 18-2
- Customizing a CLI Prompt, page 18-3
- Changing the Console Timeout Period, page 18-4

Adding a Banner

You can configure a message to display when a user connects to the security appliance, before a user logs in, or before a user enters privileged EXEC mode.

See the following guidelines:

- From a security perspective, it is important that your banner discourage unauthorized access. Do not use the words welcome or please, as they appear to invite intruders in. The following banner sets the correct tone for unauthorized access:

```
You have logged in to a secure device. If you are not authorized to access this 
device,
log out immediately or risk possible criminal consequences.
```

- See RFC 2196 for guidelines about banner messages.
- Only ASCII characters are allowed, including new line (Enter), which counts as two characters.
- Do not use tabs in the banner, because they are not preserved in the CLI version.
- There is no length limit for banners other than those for RAM and flash memory.
- You can dynamically add the hostname or domain name of the security appliance by including the strings $(hostname) and $(domain).
- If you configure a banner in the system configuration, you can use that banner text within a context by using the $(system) string in the context configuration

- After a banner is added, security appliance Telnet or SSH sessions may close if:
  - There is not enough system memory available to process the banner message(s).
  - A TCP write error occurs when attempting to display banner message(s).

To add a message of the day, login, or session banner, perform the following steps:

Step 1 From the Configuration > Device Management > Management Access > Command Line (CLI) > Banner pane, add your banner text to the field for the type of banner you are creating for the CLI:

- Session (exec) banner—This banner appears when a user accesses privileged EXEC mode at the CLI.
- Login Banner—This banner appears when a user logs in to the CLI.
- Message-of-the-day (motd) Banner—This banner appears when a user first connects to the CLI.
- ASDM Banner—This banner appears when a user connects to ASDM, following user authentication. The user is given two options for dismissing the banner:
  - Continue—Dismiss the banner and complete login as usual.
  - Disconnect— Dismiss the banner and terminate the connection.

Step 2 Click **Apply**.

The banner is added and the changes are saved to the running configuration.

Customizing a CLI Prompt

The CLI Prompt pane lets you customize the prompt used during CLI sessions. By default, the prompt shows the hostname of the security appliance. In multiple context mode, the prompt also displays the context name. You can display the following items in the CLI prompt.

context	(Multiple mode only) Displays the name of the current context.
domain	Displays the domain name.
hostname	Displays the hostname.
priority	Displays the failover priority as pri (primary) or sec (secondary).
state	Displays the traffic-passing state of the unit. The following values are displayed for the state: • act—Failover is enabled, and the unit is actively passing traffic. • stby— Failover is enabled, and the unit is not passing traffic and is in a standby, failed, or other non-active state. • actNoFailover—Failover is not enabled, and the unit is actively passing traffic. • stbyNoFailover—Failover is not enabled, and the unit is not passing traffic. This might happen when there is an interface failure above the threshold on the standby unit.

To customize the prompt used during CLI sessions so that it shows something other than the hostname or context name, complete the following steps:

Step 1 From the Configuration > Device Management > Management Access > CLI Prompt pane, do any of the following to customize the prompt:

- To add an attribute to the prompt, click the attribute in the Available Prompts list and then click **Add**. You can add multiple attributes to the prompt. The attribute is moved from the Available Prompts list to the Selected Prompts list.
- To remove an attribute from the prompt, click the attribute in the Selected Prompts list and then click **Delete**. The attribute is moved from the Selected Prompts list to the Available Prompts list.
- To change the order in which the attributes appear in the command prompt, click the attribute in the Selected Prompts list and click **Move Up** or **Move Down** to change the order.

The prompt is changed and displays in the CLI Prompt Preview field.

Step 2 Click **Apply**.

The new prompt is saved to the running configuration.

Changing the Console Timeout Period

To change the console timeout period, or the duration of time the management console remains active before automatically shutting down, perform the following steps:

Step 1 From the Configuration > Device Management > Management Access > Command Line (CLI) > Console Timeout pane, add a new timeout value in minutes.

To specify unlimited, enter 0. The default value is 0.

Step 2 Click **Apply**.

The console timeout is changed, and the changes are saved to the running configuration.

Configuring File Access

This section includes the following topics.

- Configuring the FTP Client Mode, page 18-4
- Configuring the Security Appliance as a Secure Copy Server, page 18-5
- Configuring the Security Appliance as a TFTP Client, page 18-5
- Adding Mount Points, page 18-6

Configuring the FTP Client Mode

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.

To configure the FTP client to be in passive mode, perform the following steps:

Step 1 From the Configuration > Device Management > Management Access > File Access > FTP Client pane, check **Specify FTP mode as passive**.

Step 2 Click **Apply**.

The FTP client configuration is changed and the change is saved to the running configuration.

Configuring the Security Appliance as a Secure Copy Server

You can 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.

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.

To configure the security appliance as a Secure Copy (SCP) server, perform the following steps:

Step 1 From the Configuration > Device Management > Management Access > File Access > **Secure Copy (SCP) Server** pane, check **Enable secure copy server**.

Step 2 Click **Apply**.

The changes are saved to the running configuration. The security appliance can function as an SCP server for transferring files from/to the device.

Configuring the Security Appliance as a TFTP Client

TFTP is a simple client/server file transfer protocol described in RFC783 and RFC1350 Rev. 2. You can 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 Client or Tools > Command Line Interface. In this way, you can back up and propagate configuration files to multiple security appliances.

The security appliance supports only one TFTP client. The full path to the TFTP client is specified in Configuration > Device Management > Management Access > File Access > TFTP Client. 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 client is done apart from this function.

To configure the security appliance as a TFTP client for saving configuration files to a TFTP server, perform the following steps:

Step 1 From the Configuration > Device Management > Management Access > File Access > TFTP Client pane, check **Enable**.

Step 2 From the Interface Name drop-down list, choose the interface to use as a TFTP client.

Step 3 In the IP Address field, add the IP address of the TFTP server where configuration files will be saved.

Step 4 In the Path field, add the path to the TFTP server where configuration files will be saved.

For example: /tftpboot/asa/config3

Step 5 Click **Apply**.

The changes are saved to the running configuration. This TFTP server will be used to save the security appliance configuration files. For more information, see Save the Running Configuration to a TFTP Server, page 5-4.

Adding Mount Points

Common Internet File System (CIFS) and File Transfer Protocol (FTP) mount points

This section includes the following topics:

- Adding a CIFS Mount Point, page 18-6
- Adding an FTP Mount Point, page 18-7

Adding a CIFS Mount Point

To define a CIFS mount point, perform the following steps:

Step 1 From the Configuration > Device Management > Management Access > File Access > Mount-Points pane, click **Add > CIFS Mount Point**.

The Add CIFS Mount Point dialog box appears.

Step 2 Check **Enable mount point.**

This option attaches the CIFS file system on the security appliance to the UNIX file tree.

Step 3 In the Mount Point Name field, add the name of an existing CIFS location.

Step 4 In the Server Name or IP Address field, add the name or IP address of the server where the mount point is located.

Step 5 In the Share Name field, add the name of the folder on the CIFS server.

Step 6 In the NT Domain Name field, add the name of the NT Domain where the server resides.

Step 7 In the User Name field, add the name of the user authorized for file system mounting on the server.

Step 8 In the Password field, add the password for the user authorized for file system mounting on the server.

Step 9 In the Confirm Password field, add the password again.

Step 10 Click **OK**.

The Add CIFS Mount Point dialog box closes.

Step 11 Click **Apply**.

The mount point is added to the security appliance and the change is saved to the running configuration.

Adding an FTP Mount Point

Note For an FTP mount point, the FTP Server must have a UNIX directory listing style. Microsoft FTP servers have a default of MS-DOS directory listing style.

To define an FTP mount point, perform the following steps:

Step 1 From the Configuration > Device Management > Management Access > File Access > Mount-Points pane, click **Add > FTP Mount Point**.

The Add FTP Mount Point dialog box appears.

Step 2 Check the **Enable** check box.

This option attaches the FTP file system on the security appliance to the UNIX file tree.

Step 3 In the Mount Point Name field, add the name of an existing FTP location.

Step 4 In the Server Name or IP Address field, add the name or IP address of the server where the mount point is located.

Step 5 In the Mode field, click the radio button for the FTP mode (Active or Passive). When you choose Passive mode, the client initiates both the FTP control connection and data connection. The server responds with the number of its listening port for this connection.

Step 6 In the Path to Mount field, add the directory path name to the FTP file server.

Step 7 In the User Name field, add the name of the user authorized for file system mounting on the server.

Step 8 In the Password field, add the password for the user authorized for file system mounting on the server.

Step 9 In the Confirm Password field, add the password again.

Step 10 Click **OK**.

The dialog box closes.

Step 11 Click **Apply**.

The mount point is added to the security appliance and the change is saved to the running configuration.

ICMP Access

By default, you can send ICMP packets to any security appliance interface using either IPv4 or IPv6. ICMP in IPv6 functions the same as ICMP in IPv4. ICMPv6 generates error messages, such as ICMP destination unreachable messages and informational messages like ICMP echo request and reply messages. Additionally ICMP packets in IPv6 are used in the IPv6 neighbor discovery process and path MTU discovery.

By default, the security appliance does not respond to ICMP echo requests directed to a broadcast address. You can protect the security appliance from attacks by limiting the addresses of hosts and networks that are allowed to have ICMP access to the security appliance.

Note For allowing ICMP traffic *through* the security appliance, see the "Configuring Access Rules and ACLs" section on page 21-1.

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 you configure ICMP rules, 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 rule is not configured; in that case, a **permit** statement is assumed.

To configure ICMP access rules, perform the following steps:

Step 1 From the Configuration > Device Management > Management Access > ICMP pane, click **Add**.

Step 2 Choose which version of the Internet Protocol to filter by clicking the appropriate radio button:

- **Both** (filters IPv4 and IPv6 traffic)
- **IPv4** only
- **IPv6** only

Step 3 If you want to insert a rule into the ICMP table, click the rule that the new rule will precede, and click **Insert**.

The Create ICMP Rule dialog box appears in the right-hand pane.

Step 4 From the ICMP Type drop-down list, choose the type of ICMP message for this rule.

Table 18-1 lists the types of ICMP messages.

Table 18-1 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

Table 18-1 ICMP Type Literals (continued)

ICMP Type	Literal
15	information-request
16	information-reply
17	mask-request
18	mask-reply
31	conversion-error
32	mobile-redirect

Step 5 From the Interface selection list, choose the destination security appliance interface the rule is to be applied to.

Step 6 In the IP Address field, do one of the following:

- Add a specific IP address for the host or network.
- Click **Any Address** and go to Step 9.

Step 7 From the Mask drop-down list, choose the network mask.

Step 8 Click **OK**.

The dialog box closes.

Step 9 (Optional) To set ICMP unreachable message limits, set the following options. Increasing the rate limit, along with enabling the "Decrement time to live for a connection" option on the Configuration > Firewall > Service Policy Rules > Rule Actions > Connection Settings dialog box, is required to allow a traceroute through the security appliance that shows the security appliance as one of the hops.

- Rate Limit—Sets the rate limit of unreachable messages, between 1 and 100 messages per second. The default is 1 message per second.
- Burst Size—Sets the burst rate, between 1 and 10. This keyword is not currently used by the system, so you can choose any value.

Step 10 Click **Apply**.

The ICMP rule is added to the end of the ICMP table, and the change is saved to the running configuration.

Configuring a Management Interface

A high-security interface can be identified to manage the security appliance. When a management interface is assigned, ASDM can run on it with a fixed IP address over an IPSec VPN tunnel. This is possible if VPN is configured on the security appliance and the external interface is using a dynamically assigned IP address. The management interface is also used when accessing and managing the security appliance securely from home using the VPN client.

Note If your security appliace is configured for failover, we recommend that you configure at least one management interface with the standby IP address.

To configure a management interface, perform the following steps:

Step 1 From the Configuration > Device Management > Management Access > Management Interface pane, choose the interface with the highest security (the inside interface) from the Management Access Interface drop-down list.

Step 2 Click **Apply**.

The management interface is assigned and the change is saved to the running configuration.

Configuring SNMP

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

- Information About SNMP, page 18-10
- Configuring SNMP Parameters and Management Stations, page 18-11
- Configuring SNMP Traps, page 18-14

Information About SNMP

The Simple Network Management Protocol (SNMP) enables the monitoring of network devices from a central location. The security appliance supports network monitoring using SNMP Versions 1, 2c, and 3, as well as traps and SNMP read access, but does not support SNMP write access.

You can configure the security appliance to send traps (event notifications) to a network management station (NMS), or you can use the NMS to browse the MIBs on the security appliance. Use CiscoWorks for Windows or any other SNMP MIB-II-compliant browser to receive SNMP traps and browse a MIB.

The security appliance has an SNMP agent that notifies designated management stations if events occur that are pre-defined to require a notification, for example, when a link in the network goes up or down. The notification it sends includes an SNMP OID, identifying itself to the management stations.

The security appliance SNMP agent also replies when a management station asks for information.

This section includes the following topics:

- Information About SNMP Terminology, page 18-10
- Information About the Management Information Base and Traps, page 18-11

Information About SNMP Terminology

The following terms are commonly used when working with SNMP:

Term	Description
Agent	The SNMP server running on the security appliance. The agent responds to requests for information and actions from the management station. The agent also controls access to the its management information base (MIB), the collection of objects that can be viewed or changed by the SNMP manager.
Browsing	Monitoring the health of a device from the management station by polling required information from the device SNMP agent. This activity may include issuing a series of GET-NEXT or GET-BULK requests of the MIB tree from the management station to determine values.
Management stations	The PCs or workstations set up to monitor SNMP events and manage devices such as the security appliance.
MIB	Management Information Bases, or standardized data structures, for collecting information about packets, connections, buffers, failovers, and so on. MIBs are defined by product and the protocols and hardware standards used by most network devices. SNMP management stations can browse MIBs and request specific data or events be sent as they occur. Some MIB data can be modified for administrative purposes.
OID	The system object identifier (OID) that identifies a device to its management station and indicates to users the source of information monitored and displayed.
Trap	Predefined events that generate a message from the SNMP agent to the management station. Events include alarm conditions such as linkup, linkdown, coldstart, authentication, or syslog events.

Information About the Management Information Base and Traps

MIBs are either standard or enterprise-specific. Standard MIBs are created by the IETF and documented in various RFCs. A trap reports significant events occurring on a network device, most often errors or failures. SNMP traps are defined in either standard or enterprise-specific MIBs. Standard traps are created by the IETF and documented in various RFCs. Standard traps are compiled into the security appliance software.

If needed, you can also download RFCs, standard MIBS, and standard traps from the IETF website:

http://www.ietf.org/

Download Cisco MIBs and OIDs from the following location:

http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

Download Cisco OIDs from the following location:

ftp://ftp.cisco.com/pub/mibs/oid/oid.tar.gz

Configuring SNMP Parameters and Management Stations

This section includes the following topics:

- Configuring SNMP Parameters for Versions 1 and 2c, page 18-12
- Configuring SNMP Parameters for Version 3, page 18-12
- Adding an SNMP Management Station, page 18-13

Configuring SNMP Parameters for Versions 1 and 2c

To configure SNMP parameters for Versions 1 and 2c, perform the following steps:

Step 1 In the ASDM main window, choose **Configuration** > **Device Management** > **Management Access** > **SNMP**.

Step 2 In the Community String (default) field, enter a default community string, which applies to SNMP Versions 1 and 2c only. To configure SNMP parameters for Version 3, see Configuring SNMP Parameters for Version 3, page 18-12.

Enter the password used by the SNMP management stations 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." SNMP Version 2c 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.

Step 3 In the Contact field, 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.

Step 4 In the Location field, enter the location of the security appliance being managed by SNMP. The text is case-sensitive and can be up to 127 characters. Spaces are accepted, but multiple spaces are shortened to a single space.

Step 5 In the Listening Port field, enter the number of the security appliance port that listens for SNMP requests from management stations; or keep the default, port number161.

Step 6 Click **Apply**.

SNMP parameters for Versions 1 and 2c are configured and the changes are saved to the running configuration.

Configuring SNMP Parameters for Version 3

SNMP Version 3 allows you to configure additional authentication and privacy options for more secure protocol operations by means of SNMP server groups and users. To configure SNMP parameters for Version 3, perform the following steps:

Step 1 In the ASDM main window, choose **Configuration** > **Device Management** > **Management Access** > **SNMP**.

Step 2 In the SNMPv3 Users pane, to add a configured user or a new user to a group, click **Add**. To change user parameters, click **Edit**. To remove a configured user from a group, click **Delete**. When you remove the last user in a group, ASDM deletes the group.

Note Once a user is created, you cannot change the group to which the user belongs.

The Add SNMP User Entry dialog box appears.

Step 3 In the Group Name drop-down list, choose the group to which the SNMP user will belong. The available groups are as follows:

- Auth&Encryption, in which users have authentication and encryption configured
- Authentication_Only, in which users have only authentication configured
- No_Authentication, in which users have neither authentication nor encryption configured

Step 4 In the Username field, enter the name of a configured user or a new user. The username must be unique for the SNMP server group selected.

Step 5 To have the password encrypted, click the **Encrypt Password** radio button. If you choose this option, you must enter the password as an MD5 hash value.

Step 6 Indicate the type of authentication you want to use by clicking one of the two radio buttons: **MD5** or **SHA**.

Step 7 In the Authentication Password field, type the password to use for authentication.

Step 8 Indicate the type of encryption you want to use by clicking one of these three radio buttons: **DES**, **3DES**, or **AES**.

Step 9 If you chose AES encryption, then from the AES Size drop-down list, specify which level of AES encryption to use: **128**, **192**, or **256**.

Step 10 In the Encryption Password field, type the password to use for encryption. The maximum number of characters allowed for this password is 64.

Step 11 Click **OK** to create a group (if this is the first user in that group), display this group in the Group Name drop-down list, and create a user for that group.

The Add SNMP User Entry dialog box closes.

The SNMPv3 Users pane lists the following information: SNMP Version 3 server group name, name of the user that belongs to the specified group, encrypted password setting, authentication setting, encryption algorithm setting, and the AES size setting.

Step 12 Click **Apply**.

SNMP parameters for Version 3 are configured, and the changes are saved to the running configuration.

Adding an SNMP Management Station

To add an SNMP management station, perform the following steps:

Step 1 In the ASDM main window, choose **Configuration > Device Management > Management Access > SNMP**.

Step 2 In the SNMP Management Stations pane, click **Add**.

The Add SNMP Host Access Entry dialog box appears.

Step 3 In the Interface Name drop-down list, choose the interface on which the SNMP host resides.

Step 4 In the IP Address field, enter the SNMP host IP address.

Step 5 In the UDP Port field, enter the SNMP host UDP port, or keep the default, port 162.

Step 6 In the Community String field, add the SNMP host community string. If no community string is specified for a management station, the value set in the Community String (default) field on the SNMP Management Stations pane is used.

Step 7 In the SNMP Version drop-down list, choose the SNMP version used by the SNMP host.

Step 8 If you have selected SNMP Version 3 in the previous step, in the Username drop-down list, choose the name of a configured user.

Step 9 To specify the method for communicating with this management station, check the **Poll** or **Trap** check boxes.

Step 10 Click **OK**.

The Add SNMP Host Access Entry dialog box closes.

Step 11 Click **Apply**.

The management station is configured and changes are saved to the running configuration.

Configuring SNMP Traps

To designate which traps the SNMP agent generates and how they are collected and sent to network management stations, perform the following steps:

Step 1 In the ASDM main window, choose **Configuration > Device Management > Management Access > SNMP**.

Step 2 Click **Configure Traps**.

The SNMP Trap Configuration dialog box appears.

Step 3 Check the applicable check boxes for the SNMP events to notify through SNMP traps.

Step 4 Click **OK**.

The SNMP Trap Configuration dialog box closes.

Step 5 Click **Apply**.

The SNMP traps are configured and the changes are saved to the running configuration.

Configuring Management Access Rules

Access Rules specifically permit or deny traffic to or from a particular peer (or peers), while Management Access Rules provide access control for to-the-box traffic. For example, in addition to detecting IKE Denial of Service attacks, you can block them using management access rules.

To add a Management Access Rule, perform the following steps:

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

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

- **Add Management Access Rule**
- **Add IPv6 Management Access Rule**

The appropriate Add Management Access Rule dialog box appears.

Step 3 From the Interface drop-down list, choose an interface on which to apply the rule.

Step 4 In the Action field, click one of the following:

- **Permit** (permits this traffic)
- **Deny** (denies this traffic)

Step 5 In the Source field, choose Any, or click the ellipsis (...) to browse for an address.

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

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

Step 8 (Optional) If you want to receive log messages for this access rule, check **Enable Logging**, and then from the Logging Level drop-down list, choose the log level to apply. The default level is Informational.

Step 9 (Optional) To configure advanced options, click **More Options** to configure the following settings:

- If you want to turn off this Management Access Rule, uncheck **Enable Rule**.
- Add a source service in the Source Service field, or click the ellipsis (...) to browse for a service.

  The destination service and source service must be the same. Copy and paste the destination Service field to the Source Service field.
- To configure the logging interval (if you enable logging and choose a non-default setting), enter a value in seconds in the Logging Interval field.
- To select a predefined time range for this rule, from the Time Range drop-down list, choose a time range; or click the ellipsis (...) to browse for a time range.

  The Add Time Range dialog box appears. For information about adding a time range, see Configuring Time Ranges, page 20-15.

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

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

Note After you create management access rules, you can click the radio buttons at the bottom of the pane to sort the display and show both IPv4 and IPv6 rules, IPv4 only, or IPv6 only.

Configuring AAA for System Administrators

This section describes how to enable authentication and command authorization for system administrators. Before you configure AAA for system administrators, first configure the local database or AAA server according to the "AAA Server and Local Database Support" section on page 16-3 or the "Configuring AAA Server Groups" section on page 16-9.

This section includes the following topics:

Configuring Authentication for CLI, ASDM, and enable command Access

If you enable CLI authentication, the security appliance prompts you for your username and password to log in. After you enter your information, you have access to user EXEC mode.

To enter privileged EXEC mode, enter the **enable** command or the **login** command (if you are using the local database only).

If you configure **enable** authentication, the security appliance prompts you for your username and password. If you do not configure **enable** authentication, enter the system enable password when you enter the **enable** command (set by the **enable password** command). However, if you do not use **enable** authentication, after you enter the **enable** command, you are no longer logged in as a particular user. To maintain your username, use **enable** authentication.

For authentication using the local database, you can use the **login** command, which maintains the username but requires no configuration to turn on authentication.

Note Before the security appliance can authenticate a Telnet, SSH, or HTTP user, you must first configure access to the security appliance according to the "Configuring Device Access for ASDM, Telnet, or SSH" section on page 18-1. These panes identify the IP addresses that are allowed to communicate with the security appliance.

To configure CLI, ASDM, or **enable** authentication, perform the following steps:

Step 1 To authenticate users who use the **enable** command, go to Configuration > Device Management > Users/AAA > AAA Access > Authentication, and configure the following settings:

a. Check the **Enable** check box.

b. From the Server Group drop-down list, choose a server group name or the LOCAL database.

c. (Optional) If you chose a AAA server, you can configure the security appliance to use the local database as a fallback method if the AAA server is unavailable. Click the **Use LOCAL when server group fails** check box. We recommend that you use the same username and password in the local database as the AAA server because the security appliance prompt does not give any indication which method is being used.

Step 2 To authenticate users who access the CLI or ASDM, go to Configuration > Device Management > Users/AAA > AAA Access > Authentication, and configure the following settings:

a. Check one or more of the following check boxes:

- **HTTP/ASDM**—Authenticates the ASDM client that accesses the security appliance using HTTPS. You only need to configure HTTP authentication if you want to use a AAA server. By default, ASDM uses the local database for authentication even if you do not configure this command.
- **Serial**—Authenticates users who access the security appliance using the console port.
- **SSH**—Authenticates users who access the security appliance using SSH.
- **Telnet**—Authenticates users who access the security appliance using Telnet.

b. For each service that you checked, from the Server Group drop-down list, choose a server group name or the LOCAL database.

c. (Optional) If you chose a AAA server, you can configure the security appliance to use the local database as a fallback method if the AAA server is unavailable. Click the **Use LOCAL when server group fails** check box. We recommend that you use the same username and password in the local database as the AAA server because the security appliance prompt does not give any indication which method is being used.

Step 3 Click **Apply**.

Limiting User CLI and ASDM Access with Management Authorization

If you configure CLI or **enable** authentication, you can limit a local user, RADIUS, TACACS+, or LDAP user (if you map LDAP attributes to RADIUS attributes) from accessing the CLI, ASDM, or the **enable** command.

Note Serial access is not included in management authorization, so if you enable the Authentication > Serial option, then any user who authenticates can access the console port.

To configure management authorization, perform the following steps:

Step 1 To enable management authorization, go to Configuration > Device Management > Users/AAA > AAA Access > Authorization, and check the **Perform authorization for exec shell access > Enable** check box.

This option also enables support of administrative user privilege levels from RADIUS, which can be used in conjunction with local command privilege levels for command authorization. See the "Configuring Local Command Authorization" section on page 18-20 for more information.

Step 2 To configure the user for management authorization, see the following requirements for each AAA server type or local user:

- RADIUS or LDAP (mapped) users—Configure the Service-Type attribute for one of the following values.
- RADIUS or LDAP (mapped) users—Use the IETF RADIUS numeric Service-Type attribute which maps to one of the following values.
  - Service-Type 6 (Administrative)—Allows full access to any services specified by the Authentication tab options

- Service-Type 7 (NAS prompt)—Allows access to the CLI when you configure the Telnet or SSH authentication options, but denies ASDM configuration access if you configure the HTTP option. ASDM monitoring access is allowed. If you configure **enable** authentication with the Enable option, the user cannot access privileged EXEC mode using the **enable** command.
- Service-Type 5 (Outbound)—Denies management access. The user cannot use any services specified by the Authentication tab options (excluding the Serial option; serial access is allowed). Remote-access (IPSec and SSL) users can still authenticate and terminate their remote-access sessions.

- TACACS+ users—Authorization is requested with the "service=shell" and the server responds with PASS or FAIL.
  - PASS, privilege level 1—Allows full access to any services specified by the Authentication tab options.
  - PASS, privilege level 2 and higher—Allows access to the CLI when you configure the Telnet or SSH authentication options, but denies ASDM configuration access if you configure the HTTP option. ASDM monitoring access is allowed. If you configure **enable** authentication with the Enable option, the user cannot access privileged EXEC mode using the **enable** command.
  - FAIL—Denies management access. The user cannot use any services specified by the Authentication tab options (excluding the Serial option; serial access is allowed).
- Local users—Configure the Access Restriction option. See the "Adding a User Account" section on page 16-18. By default, the access restriction is Full Access, which allows full access to any services specified by the Authentication tab options.

Configuring Command Authorization

If you want to control the access to commands, the security appliance lets you configure command authorization, where you can determine which commands that are available to a user. By default when you log in, you can access user EXEC mode, which offers only minimal commands. When you enter the **enable** command (or the **login** command when you use the local database), you can access privileged EXEC mode and advanced commands, including configuration commands.

This section includes the following topics:

- Command Authorization Overview, page 18-18
- Configuring Local Command Authorization, page 18-20
- Configuring TACACS+ Command Authorization, page 18-23

Command Authorization Overview

This section describes command authorization, and includes the following topics:

- Supported Command Authorization Methods, page 18-19
- About Preserving User Credentials, page 18-19
- Security Contexts and Command Authorization, page 18-20

Supported Command Authorization Methods

You can use one of two command authorization methods:

- Local privilege levels—Configure the command privilege levels on the security appliance. When a local, RADIUS, or LDAP (if you map LDAP attributes to RADIUS attributes) user authenticates for CLI access, the security appliance places that user in the privilege level that is defined by the local database, RADIUS, or LDAP server. The user can access commands at the user's privilege level and below. Note that all users access user EXEC mode when they first log in (commands at level 0 or 1). The user needs to authenticate again with the **enable** command to access privileged EXEC mode (commands at level 2 or higher), or they can log in with the **login** command (local database only).

Note You can use local command authorization without any users in the local database and without CLI or **enable** authentication. Instead, when you enter the **enable** command, you enter the system enable password, and the security appliance places you in level 15. You can then create enable passwords for every level, so that when you enter **enable** *n* (2 to 15), the security appliance places you in level *n*. These levels are not used unless you turn on local command authorization (see "Configuring Local Command Authorization" below). (See the *Cisco ASA 5500 Series Command Reference* for more information about **enable**.)

- TACACS+ server privilege levels—On the TACACS+ server, configure the commands that a user or group can use after they authenticate for CLI access. Every command that a user enters at the CLI is checked with the TACACS+ server.

About Preserving User Credentials

When a user logs into the security appliance, they are required to provide a username and password for authentication. The security appliance retains these session credentials in case further authentication is needed later in the session.

When the following configurations are in place, a user needs only to authenticate with the local server upon login. Subsequent serial authorization uses the saved credentials. The user is also prompted for the privilege level 15 password. When exiting privileged mode, the user is authenticated again. User credentials are not retained in privileged mode.

- Local server is configured to authenticate user access.
- Privilege level 15 command access is configured to require a password.
- User's account is configured for serial only authorization (no access to console or ASDM).
- User's account is configured for privilege level 15 command access.

The following table shows how credentials are used in this case by the security appliance.

Credentials required	Username and Password Authentication	Serial Authorization	Privileged Mode Command Authorization	Privileged Mode Exit Authorization
Username	Yes	No	No	Yes
Password	Yes	No	No	Yes
Privileged Mode Password	No	No	Yes	No

Security Contexts and Command Authorization

The following are important points to consider when implementing command authorization with multiple security contexts:

- AAA settings are discrete per context, not shared between contexts.

  When configuring command authorization, you must configure each security context separately. This provides you the opportunity to enforce different command authorizations for different security contexts.

  When switching between security contexts, administrators should be aware that the commands permitted for the username specified when they login may be different in the new context session or that command authorization may not be configured at all in the new context. Failure to understand that command authorizations may differ between security contexts could confuse an administrator. This behavior is further complicated by the next point.

- New context sessions started with the **changeto** command always use the default "enable_15" username as the administrator identity, regardless of what username was used in the previous context session. This behavior can lead to confusion if command authorization is not configured for the enable_15 user or if authorizations are different for the enable_15 user than for the user in the previous context session.

  This behavior also affects command accounting, which is useful only if you can accurately associate each command that is issued with a particular administrator. Because all administrators with permission to use the **changeto** command can use the enable_15 username in other contexts, command accounting records may not readily identify who was logged in as the enable_15 username. If you use different accounting servers for each context, tracking who was using the enable_15 username requires correlating the data from several servers.

  When configuring command authorization, consider the following:

  - An administrator with permission to use the **changeto** command effectively has permission to use all commands permitted to the enable_15 user in each of the other contexts.
  - If you intend to authorize commands differently per context, ensure that in each context the enable_15 username is denied use of commands that are also denied to administrators who are permitted use of the **changeto** command.

  When switching between security contexts, administrators can exit privileged EXEC mode and enter the **enable** command again to use the username they need.

Note The system execution space does not support AAA commands; therefore, command authorization is not available in the system execution space.

Configuring Local Command Authorization

Local command authorization lets you assign commands to one of 16 privilege levels (0 to 15). By default, each command is assigned either to privilege level 0 or 15. You can define each user to be at a specific privilege level, and each user can enter any command at their privilege level or below. The security appliance supports user privilege levels defined in the local database, a RADIUS server, or an LDAP server (if you map LDAP attributes to RADIUS attributes. See the "Configuring LDAP Attribute Maps" section on page 16-22.)

This section includes the following topics:

- Local Command Authorization Prerequisites, page 18-21

- Default Command Privilege Levels, page 18-21
- Assigning Privilege Levels to Commands and Enabling Authorization, page 18-22

Local Command Authorization Prerequisites

Complete the following tasks as part of your command authorization configuration:

- Configure **enable** authentication. (See the "Configuring Authentication for CLI, ASDM, and enable command Access" section on page 18-16.)

  enable authentication is essential to maintain the username after the user accesses the **enable** command.

  Alternatively, you can use the **login** command (which is the same as the **enable** command with authentication; for the local database only), which requires no configuration. We do not recommend this option because it is not as secure as **enable** authentication.

  You can also use CLI authentication, but it is not required.

- See the following prerequisites for each user type:
  - Local database users—Configure each user in the local database at a privilege level from 0 to 15.

    To configure the local database, see the "AAA Server and Local Database Support" section on page 16-3.

  - RADIUS users—Configure the user with Cisco VSA CVPN3000-Privilege-Level with a value between 0 and 15.
  - LDAP users—Configure the user with a privilege level between 0 and 15, and then map the LDAP attribute to Cisco VAS CVPN3000-Privilege-Level according to the "Configuring LDAP Attribute Maps" section on page 16-22.

Default Command Privilege Levels

By default, the following commands are assigned to privilege level 0. All other commands are at level 15.

- **show checksum**
- **show curpriv**
- **enable**
- **help**
- **show history**
- **login**
- **logout**
- **pager**
- **show pager**
- **clear pager**
- **quit**
- **show version**

If you move any configure mode commands to a lower level than 15, be sure to move the **configure** command to that level as well, otherwise, the user will not be able to enter configuration mode.

Assigning Privilege Levels to Commands and Enabling Authorization

To assign a command to a new privilege level, and enable authorization, follow these steps:

Step 1 To enable command authorization, go to Configuration > Device Management > Users/AAA > AAA Access > Authorization, and check **Enable authorization for command access > Enable**.

Step 2 From the Server Group drop-down list, choose **LOCAL**.

Step 3 When you enable local command authorization, you have the option of manually assigning privilege levels to individual commands or groups of commands or enabling the predefined user account privileges.

- To use predefined user account privileges, click **Set ASDM Defined User Roles**.

  The ASDM Defined User Roles Setup dialog box shows the commands and their levels. Click **Yes** to use the predefined user account privileges: Admin (privilege level 15, with full access to all CLI commands; Read Only (privilege level 5, with read-only access); and Monitor Only (privilege level 3, with access to the Monitoring section only).

- To manually configure command levels, click **Configure Command Privileges**.

  The Command Privileges Setup dialog box appears. You can view all commands by choosing **--All Modes--** from the Command Mode drop-down list, or you can choose a configuration mode to view the commands available in that mode. For example, if you choose **context**, you can view all commands available in context configuration mode. If a command can be entered in user EXEC/privileged EXEC mode as well as configuration mode, and the command performs different actions in each mode, you can set the privilege level for these modes separately.

  The Variant column displays show, clear, or cmd. You can set the privilege only for the show, clear, or configure form of the command. The configure form of the command is typically the form that causes a configuration change, either as the unmodified command (without the **show** or **clear** prefix) or as the **no** form.

  To change the level of a command, double-click it or click **Edit**. You can set the level between 0 and 15. You can only configure the privilege level of the *main* command. For example, you can configure the level of all **aaa** commands, but not the level of the **aaa authentication** command and the **aaa authorization** command separately.

  To change the level of all shown commands, click **Select All** and then **Edit**.

  Click **OK** to accept your changes.

Step 4 To support administrative user privilege levels from RADIUS, check **Perform authorization for exec shell access > Enable**.

Without this option, the security appliance only supports privilege levels for local database users and defaults all other types of users to level 15.

This option also enables management authorization for local, RADIUS, LDAP (mapped), and TACACS+ users. See the "Limiting User CLI and ASDM Access with Management Authorization" section on page 18-17 for more information.

Step 5 Click **Apply**.

Configuring TACACS+ Command Authorization

If you enable TACACS+ command authorization, and a user enters a command at the CLI, the security appliance sends the command and username to the TACACS+ server to determine if the command is authorized.

When configuring command authorization with a TACACS+ server, do not save your configuration until you are sure it works the way you want. If you get locked out because of a mistake, you can usually recover access by restarting the security appliance. If you still get locked out, see the "Recovering from a Lockout" section on page 18-27.

Be sure that your TACACS+ system is completely stable and reliable. The necessary level of reliability typically requires that you have a fully redundant TACACS+ server system and fully redundant connectivity to the security appliance. For example, in your TACACS+ server pool, include one server connected to interface 1, and another to interface 2. You can also configure local command authorization as a fallback method if the TACACS+ server is unavailable. In this case, you need to configure local users and command privilege levels according to the "Configuring Command Authorization" section on page 18-18.

This section includes the following topics:

- TACACS+ Command Authorization Prerequisites, page 18-23
- Configuring Commands on the TACACS+ Server, page 18-23
- Enabling TACACS+ Command Authorization, page 18-26

TACACS+ Command Authorization Prerequisites

Configure CLI and **enable** authentication (see the "Configuring Authentication for CLI, ASDM, and enable command Access" section on page 18-16).

Configuring Commands on the TACACS+ Server

You can configure commands on a Cisco Secure Access Control Server (ACS) TACACS+ server as a shared profile component, for a group, or for individual users. For third-party TACACS+ servers, see your server documentation for more information about command authorization support.

See the following guidelines for configuring commands in Cisco Secure ACS Version 3.1; many of these guidelines also apply to third-party servers:

- The security appliance sends the commands to be authorized as "shell" commands, so configure the commands on the TACACS+ server as shell commands.

Note Cisco Secure ACS might include a command type called "pix-shell." Do not use this type for security appliance command authorization.

- The first word of the command is considered to be the main command. All additional words are considered to be arguments, which need to be preceded by **permit** or **deny**.

  For example, to allow the **show running-configuration aaa-server** command, add **show running-configuration** to the command box, and type **permit aaa-server** in the arguments box.

- You can permit all arguments of a command that you do not explicitly deny by selecting the **Permit Unmatched Args** check box.

For example, you can configure just the **show** command, and then all the **show** commands are allowed. We recommend using this method so that you do not have to anticipate every variant of a command, including abbreviations and **?**, which show CLI usage (see Figure 18-1).

Figure 18-1 Permitting All Related Commands

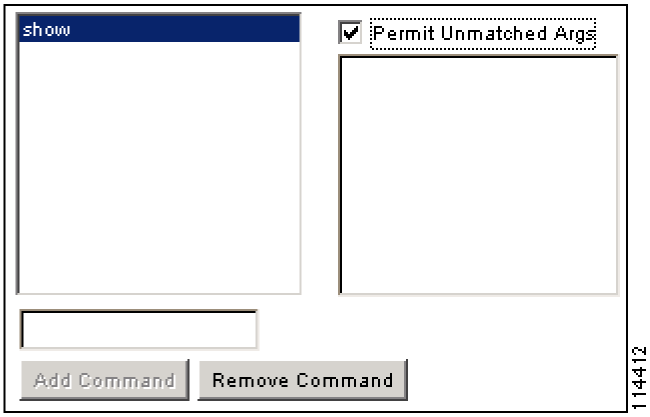
show
☑ Permit Unmatched Args

- For commands that are a single word, you *must* permit unmatched arguments, even if there are no arguments for the command, for example **enable** or **help** (see Figure 18-2).

Figure 18-2 Permitting Single Word Commands

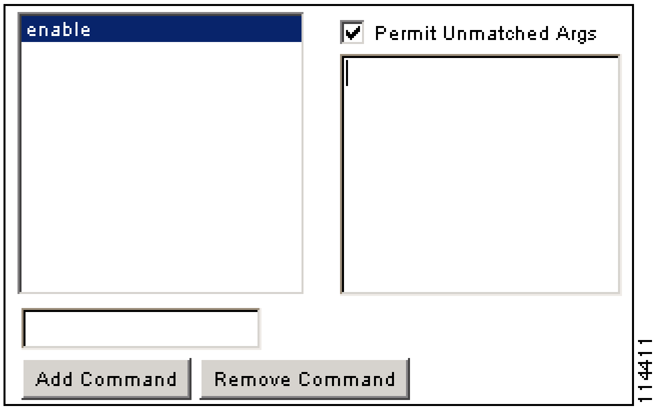
enable
☑ Permit Unmatched Args

- To disallow some arguments, enter the arguments preceded by **deny**.

  For example, to allow **enable**, but not **enable password**, enter **enable** in the commands box, and **deny password** in the arguments box. Be sure to select the Permit Unmatched Args check box so that **enable** alone is still allowed (see Figure 18-3).

Figure 18-3 Disallowing Arguments

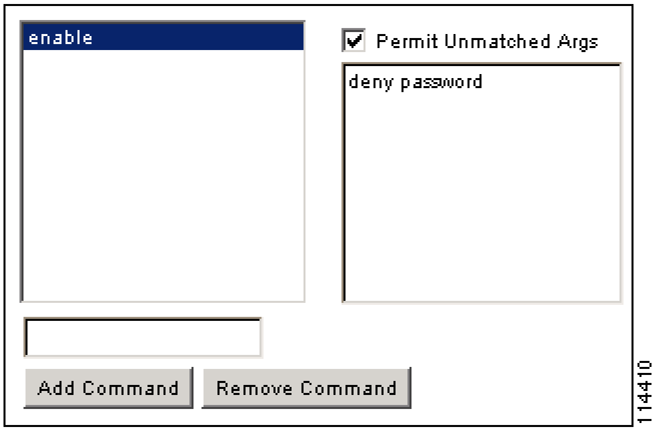

enable

☑ Permit Unmatched Args

deny password

- When you abbreviate a command at the command line, the security appliance expands the prefix and main command to the full text, but it sends additional arguments to the TACACS+ server as you enter them.

  For example, if you enter **sh log**, then the security appliance sends the entire command to the TACACS+ server, **show logging**. However, if you enter **sh log mess**, then the security appliance sends **show logging mess** to the TACACS+ server, and not the expanded command **show logging message**. You can configure multiple spellings of the same argument to anticipate abbreviations (see Figure 18-4).

Figure 18-4 Specifying Abbreviations

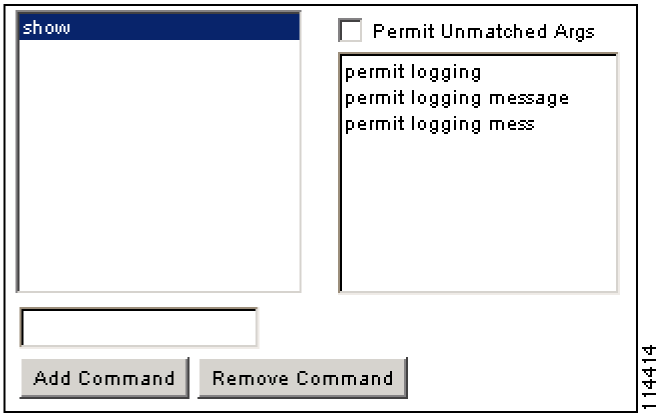

show

□ Permit Unmatched Args

permit logging
permit logging message
permit logging mess

- We recommend that you allow the following basic commands for all users:
  - **show checksum**
  - **show curpriv**
  - **enable**
  - **help**
  - **show history**
  - **login**
  - **logout**
  - **pager**

- **show pager**
- **clear pager**
- **quit**
- **show version**

Enabling TACACS+ Command Authorization

Before you enable TACACS+ command authorization, be sure that you are logged into the security appliance as a user that is defined on the TACACS+ server, and that you have the necessary command authorization to continue configuring the security appliance. For example, you should log in as an admin user with all commands authorized. Otherwise, you could become unintentionally locked out.

To configure TACACS+ command authorization, perform the following steps:

Step 1 To perform command authorization using a TACACS+ server, go to Configuration > Device Management > Users/AAA > AAA Access > Authorization, and check the **Enable authorization for command access > Enable** check box.

Step 2 From the Server Group drop-down list, choose a AAA server group name.

Step 3 (Optional) you can configure the security appliance to use the local database as a fallback method if the AAA server is unavailable. Click the **Use LOCAL when server group fails** check box. We recommend that you use the same username and password in the local database as the AAA server because the security appliance prompt does not give any indication which method is being used.

Step 4 Click **Apply**.

Configuring Management Access Accounting

To enable accounting for management access, perform the following steps:

Step 1 You can only account for users that first authenticate with the security appliance, so configure authentication using the "Configuring Authentication for CLI, ASDM, and enable command Access" section on page 18-16.

Step 2 To enable accounting of users when they enter the **enable** command:

a. Go to Configuration > Device Management > Users/AAA > AAA Access > Accounting, and check the **Require accounting to allow accounting of user activity > Enable** check box.

b. From the Server Group drop-down list, choose a RADIUS or TACACS+ server group name.

Step 3 To enable accounting of users when they access the security appliance using Telnet, SSH, or the serial console:

a. Under the Require accounting for the following types of connections area, check the check boxes for Serial, SSH, and/or Telnet.

b. For each connection type, from the Server Group drop-down list, choose a RADIUS or TACACS+ server group name.

Step 4 To configure command accounting:

a. Under the Require command accounting area, check **Enable**.

b. From the Server Group drop-down list, choose a TACACS+ server group name. RADIUS is not supported.

You can send accounting messages to the TACACS+ accounting server when you enter any command other than **show** commands at the CLI.

c. If you customize the command privilege level using the Command Privilege Setup dialog box (see the "Assigning Privilege Levels to Commands and Enabling Authorization" section on page 18-22), you can limit which commands the security appliance accounts for by specifying a minimum privilege level in the Privilege level drop-down list. The security appliance does not account for commands that are below the minimum privilege level.

Step 5 Click **Apply**.

Recovering from a Lockout

In some circumstances, when you turn on command authorization or CLI authentication, you can be locked out of the security appliance CLI. You can usually recover access by restarting the security appliance. However, if you already saved your configuration, you might be locked out. Table 18-2 lists the common lockout conditions and how you might recover from them.

Table 18-2 CLI Authentication and Command Authorization Lockout Scenarios

Feature	Lockout Condition	Description	Workaround: Single Mode	Workaround: Multiple Mode
Local CLI authentication	No users in the local database	If you have no users in the local database, you cannot log in, and you cannot add any users.	Log in and reset the passwords and **aaa** commands.	Session into the security appliance from the switch. From the system execution space, you can change to the context and add a user.
TACACS+ command authorization TACACS+ CLI authentication RADIUS CLI authentication	Server down or unreachable and you do not have the fallback method configured	If the server is unreachable, then you cannot log in or enter any commands.	1. Log in and reset the passwords and AAA commands. 2. Configure the local database as a fallback method so you do not get locked out when the server is down.	1. If the server is unreachable because the network configuration is incorrect on the security appliance, session into the security appliance from the switch. From the system execution space, you can change to the context and reconfigure your network settings. 2. Configure the local database as a fallback method so you do not get locked out when the server is down.

Table 18-2 CLI Authentication and Command Authorization Lockout Scenarios (continued)

Feature	Lockout Condition	Description	Workaround: Single Mode	Workaround: Multiple Mode
TACACS+ command authorization	You are logged in as a user without enough privileges or as a user that does not exist	You enable command authorization, but then find that the user cannot enter any more commands.	Fix the TACACS+ server user account. If you do not have access to the TACACS+ server and you need to configure the security appliance immediately, then log into the maintenance partition and reset the passwords and **aaa** commands.	Session into the security appliance from the switch. From the system execution space, you can change to the context and complete the configuration changes. You can also disable command authorization until you fix the TACACS+ configuration.
Local command authorization	You are logged in as a user without enough privileges	You enable command authorization, but then find that the user cannot enter any more commands.	Log in and reset the passwords and **aaa** commands.	Session into the security appliance from the switch. From the system execution space, you can change to the context and change the user level.

CHAPTER 19

Configuring Logging

The logging feature lets you enable logging and specify how log information is handled. The Log viewing feature lets you view syslog messages in real-time. For a description of the log viewing feature, see Chapter 46, "Monitoring Logging."

About Logging

The security appliance supports the generation of an audit trail of syslog messages that describes its activities (for example, what types of network traffic has been allowed and denied) and enables you to configure system logging.

All syslog 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 syslog messages, see the *Cisco ASA 5500 Series 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 that you view in your session are only those that are related to the current context.

Syslog 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 19-6.

Using Logging

After you have defined the security context, choose **Configuration** > **Device Management** > **Logging**. Under Logging, you can do the following:

- In the Logging Setup pane, enable logging and configure the logging parameters. For more information, see Logging Setup, page 19-2.
- In the Syslog Setup pane, set the facility code to be included in syslog 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 disable messages. For more information, see Syslog Setup, page 19-4.
- In the E-Mail Setup pane, specify syslog messages to be sent by e-mail for notification purposes. For more information, see Syslog Setup, page 19-4.
- 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 19-8.
- 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 19-10.
- 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 19-14.
- In the Syslog Server pane, specify one or more syslog servers to which the security appliance sends syslog messages. For more information, see Syslog Servers, page 19-16.
- In the SMTP pane, specify one or more SMTP servers to which the ASDM sends e-mail alerts and notification messages. For more information, see SMTP, page 19-18.
- In the NetFlow pane, export the information about the progression of a flow of packets. For more information, see Using NetFlow, page 19-18.

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. This pane also lets you change default settings for the internal log buffer and the security appliance logging queue. To access this pane, choose **Configuration** > **Device Management** > **Logging** > **Logging Setup**.

To configure logging, perform the following steps:

Step 1 Check the **Enable logging** check box to turn on logging for the main security appliance.

Step 2 Check the **Enable logging on the failover standby unit** check box to turn on logging for the standby security appliance, if available.

Step 3 Check the **Send debug messages as syslogs** check box to redirects all debugging trace output to system logs. The syslog message does not appear on the console if this option is enabled. Therefore, to view debugging messages, you must have logging enabled at the console and have it configured as the destination for the debugging syslog message number and severity level. The syslog message number to use is **711001**. The default severity level for this syslog message is debugging.

Step 4 Check the **Send syslogs in EMBLEM format** check box to enable EMBLEM format so that it is used for all log destinations, except syslog servers.

Step 5 In the Buffer Size field, specify the size of the internal log buffer to which syslog messages are saved if the logging buffer is enabled. When the buffer fills up, messages will be overwritten unless you save the logs to an FTP server or to internal flash memory. The default buffer size is 4096 bytes. The range is 4096 to 1048576.

Step 6 To save the buffer content to the FTP server before it is overwritten, check the **Save Buffer To FTP Server** check box. To allow overwriting of the buffer content, uncheck this check box.

Step 7 Click **Configure FTP Settings** to identify the FTP server and configure the FTP parameters used to save the buffer content. For more information, see Configure FTP Settings, page 19-3.

Step 8 To save the buffer content to internal flash memory before it is overwritten, check the **Save Buffer To Flash** check box.

Note This option is only available in routed or transparent single mode.

Step 9 Click **Configure Flash Usage** to specify 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 on which messages are stored. For more information, see Configure Logging Flash Usage, page 19-4.

Note This option is only available in single, routed or transparent mode.

Step 10 In the Queue Size field, specify the queue size for system logs that are to be viewed in 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
•	•	•	•	—

Configure FTP Settings

The Configure FTP Settings dialog box lets you specify the configuration for the FTP server that is used to save the log buffer content.

To configure FTP settings, perform the following steps:

Step 1 Check the **Enable FTP client** check box to enable configuration of the FTP client.

Step 2 In the Server IP Address field, specify the IP address of the FTP server.

Step 3 In the Path field, specify the directory path on the FTP server to store the saved log buffer content.

Step 4 In the Username field, specify the username to log in to the FTP server.

Step 5 In the Password field, specify the password associated with the username to log in to the FTP server.

Step 6 In the Confirm Password field, reenter the password, and click **OK**.

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

The Configure Logging Flash Usage dialog box lets you specify the limits for saving log buffer content to internal flash memory.

To configure logging flash usage, perform the following steps:

Step 1 In the Maximum Flash to Be Used by Logging field, specify the maximum amount of internal flash memory that can be used for logging (in KB).

Step 2 In the Minimum Free Space to Be Preserved field, specify the amount of internal flash memory that is preserved (in KB). When the internal flash memory approaches that limit, new logs are no longer saved.

Step 3 Click **OK** to close this 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
•	•	•	—	—

Syslog Setup

The Syslog Setup pane lets you set the facility code to include in messages destined for syslog servers and determine whether syslog messages should include the timestamp. You can change message severity levels and disable messages that you do not want to be logged. To access this pane, choose **Configuration > Device Management > Logging > Syslog Setup**.

To configure syslog messaging, perform the following steps:

Step 1 From the Facility code to include in syslogs drop-down list, choose 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 eight available facilities, you might need to change this value for system logs.

Step 2 To add the date and time in each syslog message sent, check the **Include timestamp in syslogs** check box.

Step 3 From the Show drop-down list, choose the information to be displayed in the Syslog ID table. Available options are as follows:

- To specify that the Syslog ID table should display the entire list of syslog message IDs, choose **Show all syslog IDs**.
- To specify that the Syslog ID table should display only those syslog message IDs that have been explicitly disabled, choose **Show disabled syslog IDs**.
- To specify that the Syslog ID table should display only those syslog message IDs with severity levels that have changed from their default values, choose **Show syslog IDs with changed logging**.
- To specify that the Syslog ID table should display only those syslog message IDs with severity levels that have been modified and the IDs of syslog messages that have been explicitly disabled, choose **Show syslog IDs that are disabled or with a changed logging level**.

Step 4 The Syslog ID Setup Table displays the list of syslog messages based on the setting in the Syslog ID Setup Table. Choose individual messages or ranges of message IDs that you want to modify. You can either disable 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.

Step 5 To configure syslog messages to include a device ID, click **Advanced**. For more information, see Edit Syslog ID Settings, page 19-5 and Advanced Syslog Configuration, page 19-6.

Modes

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

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

Edit Syslog ID Settings

The Edit Syslog ID Settings dialog box lets you modify the severity level of the selected syslog messages or specify that the selected syslog messages should be disabled.

To change syslog message settings, perform the following steps:

Note The Syslog ID(s) field is display-only. The values that appear in this area are determined by the entries you chose in the Syslog ID table, located in the Syslog Setup pane.

Step 1 Check the **Disable Message(s)** check box to disable messages for the syslog message ID(s) displayed in the Syslog ID(s) list.

Step 2 From the Logging Level drop-down list, choose the severity level of messages to be sent for the syslog message ID(s) displayed in the Syslog ID(s) list. Severity levels are defined as follows:

- Emergency (level 0, system is unusable)
- Alert (level 1, immediate action is needed)
- Critical (level 2, critical conditions)
- Error (level 3, error conditions)
- Warning (level 4, warning conditions)
- Notification (level 5, normal but significant conditions)
- Informational (level 6, informational messages only)
- Debugging (level 7, debugging messages only)

Step 3 Click **OK** to close this 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
•	•	•	•	—

Advanced Syslog Configuration

You can configure the security appliance to include a device ID in non-EMBLEM-formatted syslog messages. You can specify only one type of device ID for syslog messages. The device ID can be the hostname of the adaptive security appliance, an interface IP address, the context, or a text string.

The Advanced Syslog Configuration dialog box lets you determine whether syslog messages should include a device ID. If this feature is enabled, the device ID is automatically included in all non-EMBLEM formatted syslog messages.

To specify additional syslog message settings, perform the following steps:

Step 1 Check the **Enable syslog device ID** check box to specify that a device ID should be included in all non-EMBLEM formatted syslog messages.

Step 2 To specify which to use as the device ID, choose one of the following options:

- Hostname
- IP address

  Choose the interface name that corresponds to the specified IP address from the drop-down list.
- String

In the User-Defined ID field, specify an alphanumeric, user-defined string.

Step 3 Click **OK** to close this 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
•	•	•	•	—

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 syslog messages to be sent as e-mail messages for notification purposes. You can filter the syslog messages sent to a destination e-mail address by severity level. The table shows which entries have been created. To access this pane, choose **Configuration** > **Device Management** > **Logging** > **E-Mail Setup**.

To configure e-mail to send notification of selected syslog messages, perform the following steps:

Step 1 In the Source E-Mail Address field, specify the e-mail address that is used as the source address for syslog messages that are sent as e-mail messages.

Step 2 Click **Add** to enter a new e-mail address recipient of the specified syslog messages.

Step 3 Choose the severity level of the syslog messages that are sent to the recipient from the drop-down list. The syslog 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. For more information, see Logging Filters, page 19-10.

Step 4 Click **Edit** to modify an existing the severity level of the syslog messages that are sent to this recipient.

Step 5 Click **OK** to close this 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/Edit E-Mail Recipients

The Add/Edit E-Mail Recipient dialog box lets you set up a destination e-mail address for a specified severity of syslog 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 specified in this dialog box and the global filter set for all e-mail recipients in the Logging Filters pane.

To add or edit e-mail recipients and severity levels, see Syslog Setup, page 19-4.

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

The Event Lists pane lets you create custom lists of events that are used to choose which syslog messages are sent to a specific destination. After you enable logging and configure the logging parameters using the Logging Setup pane, create one or more lists of events in the Event Lists pane. Use these event lists in the Logging Filters pane to specify a logging destination for each list of events. To access this pane, choose **Configuration** > **Device Management** > **Logging** > **Event Lists**.

You use three criteria to define an event list:

- Message Class
- Severity
- Message ID

A message class is a group of syslog messages related to a security appliance feature that enables you to specify an entire class of messages rather than specifying a class for each message individually. For example, use the auth class to select all syslog messages that are related to user authentication.

Severity level classifies syslog messages based on the relative importance of the event in the normal functioning of the network. The highest severity level is emergency, which means the resource is no longer available. The lowest severity level 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 syslog messages, such as 101001-1990120.

To create custom lists of events to send to a specific logging destination, perform the following steps:

Step 1 Click **Add** to display the Add Event List dialog box.

Step 2 In the Name field, enter the name of the event list. No spaces are allowed.

Step 3 In the Event Class/Severity area, click **Add** to display the Add Class and Severity Filter dialog box.

Step 4 Choose the event class from the drop-down list. Available event classes include the following:

- 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

Step 5 Choose the severity level from the drop-down list. Severity levels include the following:

- Emergency (level 0, system is unusable)
- Alert (level 1, immediate action is needed)
- Critical (level 2, critical conditions)
- Error (level 3, error conditions)
- Warning (level 4, warning conditions)
- Notification (level 5, normal but significant conditions)
- Informational (level 6, informational messages only)
- Debugging (level 7, debugging messages only)

Step 6 Click **OK** to close this dialog box.

Step 7 In the Message ID Filters area, click **Add** to display the Add Syslog Message ID Filter dialog box.

Step 8 In the Message IDs field, enter a syslog message ID or range of IDs (for example, 101001-199012) to include in the filter.

Step 9 Click **OK** to close this dialog box.

The event of interest appears in the list. To change this entry, click **Edit**.

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 Event List

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 level, or by message ID.

To add or edit an event list, see Event Lists, page 19-8.

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

The Add/Edit Syslog Message ID Filter dialog box lets you specify one or more syslog message IDs to be included in the event list.

To add or edit a syslog message ID filter, see Event Lists, page 19-8.

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

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 in the Event Lists pane. To access this pane, choose **Configuration** > **Device Management** > **Logging** > **Logging Filters**.

To apply message filters to a log destination, perform the following steps:

Step 1 Choose the name of the logging destination to which you want to apply a filter. Available logging destinations are as follows:

- Console
- Security appliance
- Syslog Servers
- SNMP Trap
- E-Mail
- Internal Buffer
- Telnet Sessions

Included in this selection are the second column, Syslogs From All Event Classes, and the third column, Syslogs From Specific Event Classes. The second column lists the severity or the event class to use to filter messages for the log destination, or whether logging is disabled for all event classes. The third column lists the event class to use to filter messages for that log destination. For more information, see Add/Edit Syslog Message ID Filter, page 19-10, Add/Edit Class and Severity Filter, page 19-13, and Event Lists, page 19-8.

Step 2 Click **Edit** to display the Edit Logging Filters dialog box. To apply, edit, or disable filters, see Edit Logging Filters, page 19-11.

Modes

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

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

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 logging from all event classes. You can filter messages according to message class and severity level, or use an event list that you create in the Event Lists pane.

The selected logging destination for this filter appears at the top.

To apply filters, perform the following steps:

Step 1 Choose the **Filter on severity** option to filter syslog messages according to their severity level.

Step 2 Choose the **Use event list** option to filter syslog messages according to an event list.

Step 3 Choose the **Disable logging from all event classes** option to disable all logging to the selected destination.

Step 4 Click **New** to add a new event list. To add a new event list, see Event Lists, page 19-8.

Step 5 Choose the event class from the drop-down list. Available event classes include the following:

- 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

Step 6 Choose the level of logging messages from the drop-down list. Severity levels include the following:

- Emergency (level 0, system is unusable)
- Alert (level 1, immediate action is needed)
- Critical (level 2, critical conditions)
- Error (level 3, error conditions)
- Warning (level 4, warning conditions)
- Notification (level 5, normal but significant conditions)
- Informational (level 6, informational messages only)
- Debugging (level 7, debugging messages only)

Step 7 Click **Add** to add the event class and severity level, and then click **OK**.

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

The Add/Edit Class and Severity Filter dialog box lets you specify a message class and severity level to be used to filter messages.

To add or edit a message class and severity level for filtering messages, perform the following steps:

Step 1 Choose the event class from the drop-down list. Available event classes include the following:

- 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

Step 2 Choose the level of logging messages from the drop-down list. Severity levels include the following:

- Emergency (level 0, system is unusable)
- Alert (level 1, immediate action is needed)
- Critical (level 2, critical conditions)
- Error (level 3, error conditions)
- Warning (level 4, warning conditions)
- Notification (level 5, normal but significant conditions)
- Informational (level 6, informational messages only)
- Debugging (level 7, debugging messages only)

Step 3 Click **OK** when you are done making selections.

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

The Add/Edit Syslog Message ID Filter dialog box lets you specify individual syslog message IDs or ranges of IDs to include in the event list filter.

To add or edit a syslog message ID filter, see Event Lists, page 19-8.

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

The Rate Limit pane lets you specify the number of syslog messages that the firewall can send. You can specify a rate limit for message logging levels or 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. To access this pane, choose **Configuration > Device Management > Logging > Rate Limit**.

To assign rate limits for all syslog messages in a logging level, perform the following steps:

Step 1 Choose the logging level (message severity level) to which you want to assign rate limits. Severity levels are defined as follows:

Description	Severity Level
Disabled	No logging
Emergency	0—System is unusable
Alert	1—Immediate action is needed
Critical	2—Critical conditions
Error	3—Error conditions
Warning	4—Warning conditions
Notification	5—Normal but significant conditions
Informational	6—Informational messages only
Debugging	7—Debugging messages only

Step 2 The No of Messages field displays the number of messages sent. The Interval (Seconds) field displays the interval, in seconds, that is used to limit how many messages at this logging level can be sent. Choose a logging level from the table and click **Edit** to display the Edit Rate Limit for Syslog Logging Level dialog box. To continue, see Edit Rate Limit for Syslog Logging Level, page 19-15.

To assign or change rate limits to individual syslog messages, perform the following steps:

Step 1 To assign the rate limit of a specific syslog message, click **Add** to display the Add Rate Limit for Syslog Message dialog box. To continue, see Add/Edit Rate Limit for Syslog Message, page 19-16.

Step 2 To change the rate limit of a specific syslog message, click **Edit** to display the Edit Rate Limit for Syslog Message dialog box. To continue, see Add/Edit Rate Limit for Syslog Message, page 19-16.

Modes

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

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

Edit Rate Limit for Syslog Logging Level

The Edit Rate Limit for Syslog Logging Level **dialog** box lets you limit the number of messages that the adaptive security appliance can send in a specified time interval. The selected message severity level displays.

To change the rate limit of a specified logging level, perform the following steps:

Step 1 Enter the maximum number of messages at this logging level that can be sent.

Step 2 Enter the amount of time, in seconds, that is used to limit the rate of messages at this logging level, and click **OK**.

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

The Add/Edit Rate Limit for Syslog Message dialog box lets you assign rate limits to a specific syslog message.

To add or change the rate limit for a specific syslog message, perform the following steps:

Step 1 To add a rate limit to a specific syslog message, click **Add** to display the Add Rate Limit for Syslog Message dialog box. To change a rate limit for a syslog message, click **Edit** to display the Edit Rate Limit for Syslog Message dialog box.

Step 2 Enter the message ID of the syslog message that you want to limit.

Step 3 Enter the maximum number of messages that can be sent in the specified time interval.

Step 4 Enter the amount of time, in seconds, that is used to limit the rate of the specified message, and click **OK**.

Note To allow an unlimited number of messages, leave both the Number of Messages and Time Interval fields 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

The Syslog Servers pane lets you specify the syslog servers to which the adaptive security appliance should send syslog messages. To use the syslog server(s) you define, you must enable logging using the Logging Setup pane and set up the available destinations in the Logging Filters pane. To access this pane, choose **Configuration** > **Device Management** > **Logging** > **Syslog Server**.

To specify the syslog servers to which the adaptive security appliance should send syslog messages, perform the following steps:

Step 1 To add a new syslog server, click **Add** to display the Add Syslog Server dialog box. To change an existing syslog server settings, click **Edit** to display the Edit Syslog Server dialog box.

Step 2 Specify the number of messages that are allowed to be queued on the adaptive security appliance when a syslog server is busy. A zero value means an unlimited number of messages may be queued.

Step 3 Check the **Allow user traffic to pass when TCP syslog server is down** check box to specify whether or not to restrict all traffic if any syslog server is down.

Step 4 To continue, see Add/Edit Syslog Server, page 19-17.

Note You can set up a maximum of four syslog servers per security context (up to a total of 16).

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 Server

The Add/Edit Syslog Server dialog box lets you add or edit the syslog servers to which the adaptive security appliance sends syslog messages. To use the syslog server(s) you define, you must enable logging in the Logging Setup pane and set up the specific filters for log destinations in the Logging Filters pane.

To add or edit a syslog server, perform the following steps:

Step 1 Choose the interface used to communicate with the syslog server from the drop-down list.

Step 2 Enter the IP address that is used to communicate with the syslog server.

Step 3 Choose the protocol (either TCP or UDP) that is used by the syslog server to communicate with the security appliance.

Step 4 Enter the port number used by the syslog server to communicate with the adaptive security appliance.

Step 5 Check the **Log messages in Cisco EMBLEM format (UDP only)** check box to specify whether to log messages in Cisco EMBLEM format (available only if UDP is selected as the protocol).

Step 6 Check the **Enable secure logging using SSL/TLS (TCP only)** check box to specify that the connection to the syslog server is secure through the use of SSL/TLS over TCP, and that the syslog message content is encrypted.

Step 7 Click **OK** to complete 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
•	•	•	•	—

SMTP

The SMTP pane allows you to configure the remote SMTP server IP address to which e-mail alerts and notifications are sent in response to specific events. To access this pane, choose **Configuration** > **Device Setup** > **Logging** > **SMTP**.

To configure the remote SMTP server, perform the following steps:

Step 1 Enter the IP address of the primary SMTP server.

Step 2 (Optional) Enter the IP address of the standby SMTP server, and click **Apply**.

Modes

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

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

Using NetFlow

The NetFlow pane lets you enable the transmission of data about a flow of packets. To access this pane, choose **Configuration** > **Device Management** > **Logging** > **NetFlow**.

Note IP address and hostname assignments should be unique throughout the NetFlow configuration.

To use NetFlow, perform the following steps:

Step 1 Enter the template timeout rate, which is the interval (in minutes) at which template records are sent to all configured collectors. The default value is 30 minutes.

Step 2 To delay the export of flow-creation events and process a single flow-teardown event instead of a flow-creation event and a flow-teardown event, check the **Delay export of flow creation events for short-lived flows** check box, and then enter the number of seconds for the delay in the Delay By field.

Step 3 Specify the collector(s) to which NetFlow packets will be sent. You can configure a maximum of five collectors. To configure a collector, click **Add** to display the Add Collector dialog box, and perform the following steps:

a. Enter the IP address or hostname and the UDP port number in the associated fields.

b. Choose the interface to which NetFlow packets will be sent from the drop-down list.

Step 4 To configure more collectors, repeat Step 2 for each additional collector, and click **OK**.

Step 5 To change collector configuration details, select a collector and click **Edit**. To remove a configured collector, select it and click **Delete**.

Step 6 When NetFlow is enabled, certain syslog messages become redundant. To maintain system performance, we recommend that you disable all redundant syslog messages, because the same information is exported through NetFlow. To disable all redundant syslog messages, check the **Disable redundant syslog messages** check box. To display the redundant syslog messages and their status, click **Show Redundant Syslog Messages**.

The Redundant Syslog Messages dialog box appears. The Syslog ID field displays the redundant syslog message numbers. The Disabled field indicates whether or not the specified syslog message is disabled. Click **OK** to close this dialog box.

To disable individual redundant syslog messages, choose **Configuration** > **Device Management** > **Logging** > **Syslog Setup**.

Step 7 To continue, see the "Matching NetFlow Events to Configured Collectors" section on page 19-19.

Step 8 Click **Apply** to save your changes. Click **Reset** to enter new settings.

Matching NetFlow Events to Configured Collectors

After you configure NetFlow collectors, you can match a NetFlow event with any of these configured collectors.

To specify which NetFlow events should be sent to which collector, perform the following steps:

Step 1 In the ASDM main application window, choose **Configuration** > **Firewall** > **Service Policy Rules**.

Step 2 Choose **Global Policy** in the table, and click **Add** to display the Add Service Policy Rule dialog box. For more information about service policy rules, see the "Adding a Service Policy Rule for Through Traffic" section on page 23-7.

Note NetFlow actions are available only for global service policy rules and are applicable only to the class-default traffic class and to traffic classes with traffic match criteria of "Source and Destination IP Address (uses ACL)" or "Any traffic."

Step 3 Click the **Rule Actions** tab, and then click the **NetFlow** tab.

Step 4 Click **Add** to display the Add Flow Event dialog box.

Step 5 Choose the flow event type from the drop-down list. Available options are created, torn down, denied, or all events.

Step 6 Choose collectors to which you want events sent by checking the corresponding check boxes in the Send column.

Step 7 To add, edit or delete collectors, click **Manage** to display the list of configured collectors in the Manage NetFlow Collectors dialog box. To continue, see Step 3 of the "Using NetFlow" section on page 19-18.

Step 8 To change settings for a configured collector, select it from the list and click **Edit**. To remove a collector from this list, select it from the list and click **Delete**.

Step 9 In the Redundant Syslog Messages area, to disable redundant syslog messages and maintain current performance levels, check the **Disable redundant syslog messages** check box. Click **Show Redundant Syslog Messages** to display a list of redundant syslog messages and their status (disabled or not). You can disable or enable individual syslog messages later by choosing **Configuration** > **Device Management** > **Logging**. Click **OK** to close the Redundant Syslog Messages dialog box.

Step 10 Click **OK** to close the Manage NetFlow Collectors dialog box and return to the Add Flow Event dialog box. Click **OK** again to close the Add Flow Event dialog box and return to the NetFlow tab.

Step 11 To change flow event entries, choose an entry from the list, and click **Edit**. To remove flow event entries, choose an entry from the list, and click **Delete**.

Step 12 Click **Finish** to exit the wizard.

For more information about NetFlow, see the *Cisco ASA 5500 Series Configuration Guide using the CLI* and the *Implementation Note for NetFlow Collectors*.

Modes

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

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

PART 3

Configuring the Firewall

CHAPTER 20

Adding Global Objects

The 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 Objects pane to change it in a single place.

This chapter includes the following sections:

- Using Network Objects and Groups, page 20-1
- Configuring Service Groups, page 20-5
- Configuring Class Maps, page 20-8
- Configuring Inspect Maps, page 20-8
- Configuring Regular Expressions, page 20-8
- Configuring TCP Maps, page 20-14
- Configuring Global Pools, page 20-14
- Configuring Time Ranges, page 20-15

Using Network Objects and Groups

This section describes how to use network objects and groups, and includes the following topics:

- Network Object Overview, page 20-2
- Configuring a Network Object, page 20-2
- Configuring a Network Object Group, page 20-3
- Using Network Objects and Groups in a Rule, page 20-4
- Viewing the Usage of a Network Object or Group, page 20-5

Network Object Overview

Network objects 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 rule or a AAA rule, you can choose these predefined addresses instead of typing them in manually. Moreover, if you change the definition of an object, the change is inherited automatically by any rules using the object.

You can add network objects manually, or you can let ASDM automatically create objects from existing configuration, such as access rules and AAA rules. If you edit one of these derived objects, it persists even if you later delete the rule that used it. Otherwise, derived objects only reflect the current configuration if you refresh.

A network object group is a group containing multiple hosts and networks together. A network object group can also contain other network object groups. You can then specify the network object group as the source address or destination address in an access rule.

When you are configuring rules, the ASDM window includes an Addresses side pane at the right that shows available network objects and network object groups; you can add, edit, or delete objects directly in the Addresses pane. You can also drag additional network objects and groups from the Addresses pane to the source or destination of a selected access rule.

Configuring a Network Object

To configure a network object, perform the following steps:

Step 1 Choose **Configuration** > **Firewall** > **Objects** > **Network Objects/Group**.

Step 2 Click **Add**, and choose **Network Object** to add a new object, or choose an existing object to edit, and click **Edit**.

You can also add or edit network objects from the Addresses side pane in a rules window or when you are adding a rule.

To find an object in the list, enter a name or IP address in the Filter field and click **Filter**. The wildcard characters asterisk (*) and question mark (?) are allowed.

The Add/Edit Network Object dialog box appears.

Step 3 Fill in the following values:

- Name—(Optional) The object name. Use characters a to z, A to Z, 0 to 9, a dot, a dash, or an underscore. The name must be 64 characters or less.
- IP Address—An IPv4 or an IPv6 address, either a host or network address. When you enter a colon (:) in this field for an IPv6 address, the Netmask field changes to Prefix Length.
- Netmask or Prefix Length—If the IP address is an IPv4 address, enter the subnet mask. If the IP address is an IPv6 address, enter the prefix.
- Description—(Optional) The description of the network object.

Step 4 Click **OK**.

Step 5 Click **Apply** to save the configuration.

You can now use this network object when you create a rule. For an edited object, the change is inherited automatically by any rules using the object.

Note You cannot delete a network object that is in use.

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 Network Object Group

To configure a network object group, perform the following steps:

Step 1 In the Configuration > Firewall > Objects > Network Objects/Group pane, click **Add > Network Object Group** to add a new object group, or choose an object group and click **Edit**.

You can also add or edit network object groups from the Addresses side pane in a rules window, or when you are adding a rule.

To find an object in the list, enter a name or IP address in the Filter field and click Filter. The wildcard characters asterisk (*) and question mark (?) are allowed.

The Add/Edit Network Object Group dialog box appears.

Step 2 In the Group Name field, enter a group name.

Use characters a to z, A to Z, 0 to 9, a dot, a dash, or an underscore. The name must be 64 characters or less.

Step 3 (Optional) In the Description field, enter a description up to 200 characters in length.

Step 4 You can add existing objects or groups to the new group (nested groups are allowed), or you can create a new address to add to the group:

- To add an existing network object or group to the new group, double-click the object in the Existing Network Objects/Groups pane.

  You can also select the object, and then click **Add**. The object or group is added to the right-hand Members in Group pane.

- To add a new address, fill in the values under the Create New Network Object Member area, and click **Add**.

  The object or group is added to the right-hand Members in Group pane. This address is also added to the network object list.

To remove an object, double-click it in the Members in Group pane, or click **Remove**.

Step 5 After you add all the member objects, click OK.

You can now use this network object group when you create a rule. For an edited object group, the change is inherited automatically by any rules using the group.

Note You cannot delete a network object group that is in use.

Modes

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

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

Using Network Objects and Groups in a Rule

When you create a rule, you can enter an IP address manually, or you can browse for a network object or group to use in the rule. To use a network object or group in a rule, perform the following steps:

Step 1 From the rule dialog box, click the **...** browse button next to the source or destination address field.

The Browse Source Address or Browse Destination Address dialog box appears.

Step 2 You can either add a new network object or group, or choose an existing network object or group by double-clicking it.

To find an object in the list, enter a name or IP address in the Filter field and click **Filter**. The wildcard characters asterisk (*) and question mark (?) are allowed.

- To add a new network object, see the "Configuring a Network Object" section on page 20-2.
- To add a new network object group, see the "Configuring a Network Object Group" section on page 20-3.

After you add a new object or double-click an existing object, it appears in the Selected Source/Destination field. For access rules, you can add multiple objects and groups in the field, separated by commas.

Step 3 Click **OK**.

You return to the 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
•	•	•	•	—

Viewing the Usage of a Network Object or Group

To view what rules use a network object or group, in the Configuration > Firewall > Objects > Network Objects/Group pane, click the magnifying glass Find icon.

The Usages dialog box appears listing all the rules currently using the network object or group. This dialog box also lists any network object groups that contain the object.

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, page 20-5
- Add/Edit Service Group, page 20-6
- Browse Service Groups, page 20-7

Service Groups

The Service Groups pane lets you associate multiple services into a named group. You can specify any type of protocol and service in one group or 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 Services pane at the right that shows available service groups and other global objects; you can add, edit, or delete objects directly in the Services pane.

Fields

- Add—Adds a service group. Choose the type of service group to add from the drop-down list or choose Service Group for multiple types.
- 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.
- 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.
- Protocol—Lists the service group protocols.
- Source Ports—Lists the protocol source ports.
- Destination Ports—Lists the protocol destination ports.
- ICMP Type—Lists the service group ICMP type.
- 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

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 Add/Edit TCP Service Group dialog box is shown.

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.
- Existing Service/Service Group—Identifies items that can be added to the service group. Choose from already defined service groups, or choose from a list of commonly used port, type, or protocol names.
  - Service Groups—The title of this table depends on the type of service group you are adding. It includes the defined service groups.
  - Predefined—Lists the predefined ports, types, or protocols.
- Create new member—Lets you create a new service group member.

- Service Type—Lets you select the service type for the new service group member. Service types include TCP, UDP, TCP-UDP, ICMP, and protocol.
- Destination Port/Range—Lets you enter the destination port or range for the new TCP, UDP, or TCP-UDP service group member.
- Source Port/Range—Lets you enter the source port or range for the new TCP, UDP, or TCP-UDP service group member.
- ICMP Type—Lets you enter the ICMP type for the new ICMP service group member.
- Protocol—Lets you enter the protocol for the new protocol service group member.

- 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

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

For information about class maps, see the "Class Map Field Descriptions" section on page 25-42.

Configuring Inspect Maps

For information about inspect maps, see the "Inspect Map Field Descriptions" section on page 25-64.

Configuring Regular Expressions

This section describes how to configure regular expressions, and includes the following topics:

- Regular Expressions, page 20-8
- Add/Edit Regular Expression, page 20-9
- Build Regular Expression, page 20-11
- Test Regular Expression, page 20-13
- Add/Edit Regular Expression Class Map, page 20-14

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

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 20-1, or you can click **Build** to use the Build Regular Expression dialog box.

Note As an optimization, the security appliance searches on the deobfuscated URL. Deobfuscation compresses multiple forward slashes (/) into a single slash. For strings that commonly use double slashes, like "http://", be sure to search for "http:/" instead.

Table 20-1 lists the metacharacters that have special meanings.

Table 20-1 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_**}** or **{**_x_**,}**	Minimum repeat quantifier	Repeat at least _x_ times. For example, **ab(xy){2,}z** matches abxyxyz, abxyxyxyz, and so on.
[_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.
\	Escape character	When used with a metacharacter, matches a literal character. For example, **\[** matches the left square bracket.

Table 20-1 regex Metacharacters (continued)

Character	Description	Notes
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

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.

Note As an optimization, the security appliance searches on the deobfuscated URL. Deobfuscation compresses multiple forward slashes (/) into a single slash. For strings that commonly use double slashes, like “http://”, be sure to search for “http:/” instead.

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 automatically 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 Occurrences—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

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 the 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

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
•	•	•	•	—

Configuring TCP Maps

For information about TCP maps, see the "Enabling Connection Limits, TCP Normalization , and TCP State Bypass" section on page 28-10.

Configuring Global Pools

For information about global pools, see the "Using Dynamic NAT" section on page 22-16.

Configuring Time Ranges

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 recurring entries.

For detailed steps on adding a time range to an access rule, see the "Adding a Time Range to an Access Rule" section on page 20-15.

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.
- 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

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. For detailed steps on adding a time range to an IPv6 ACL, see the "Configuring Access Rules and ACLs" section on page 21-8.

Adding a Time Range to an Access Rule

You can add a time range to an ACL to specify when traffic can be allowed or denied through an interface.

To add a time range to an ACL, perform the following steps:

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

Step 2 Choose the access list type by clicking the **IPv4 Only**, the **IPv6 Only**, or the **IPV6 and IPv6** radio button.

Step 3 Click **Add**. The Add Access Rule window appears.

Step 4 From the Interface drop down list, choose the desired interface.

The management interface is for management only and cannot be used to configure an access rule.

Step 5 Click **Permit** or **Deny** to permit or deny the action.

Step 6 In the Source field, enter an IP address.

Step 7 In the Destination field, enter an IP address.

Step 8 Select the service type.

Step 9 Click **More Options** to expand the list.

Step 10 To the right of the Time Range drop down list, click the browse button.

The Browse Time Range window appears.

Step 11 Click **Add**.

The Add Time Range window appears.

Step 12 In the Time Range Name field, enter a time range name, with no spaces.

Step 13 Choose the Start Time and the End Time by doing one of the following:

a. Allow the default settings, in which the Start Now and the Never End radio buttons are checked.

b. Apply a specific time range by clicking the **Start at** and **End at** radio buttons and selecting the specified start and stop times from the lists.

The time range is inclusive of the times that you enter.

Step 14 (Optional) To specify additional time constraints for the time range, such as specifying the days of the week or the recurring weekly interval in which the time range will be active, click **Add**, and do one of the following:

a. Click **Specify days of the week and times on which this recurring range will be active**, and choose the days and times from the lists, and click **OK**.

b. Click **Specify a weekly interval when this recurring range will be active**, and choose the days and times from the lists, and click **OK**.

Step 15 Click OK to apply the time range.

Step 16 Click OK to apply the access rule.

Note Creating a time range does not restrict access to the device. This pane defines the time range only.

Add/Edit Time Range Field Descriptions

- 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.
- Recurring Time Ranges—Configures daily or weekly time ranges.
  - Add—Adds a recurring time range.
  - Edit—Edits the selected recurring time range.
  - Delete—Deletes the 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

The Add/Edit Recurring Time Range pane lets you fine time ranges further by letting you configure them on a daily or weekly basis.

For detailed steps on adding a recurring time range to an access rule, see the “Adding a Time Range to an Access Rule” section on page 20-15.

Note Creating a time range does not restrict access to the device. This pane defines the time range only.

Add/Edit Recurring Time Range Field Descriptions

- 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
•	•	•	•	—

CHAPTER 21

Configuring Access Rules and ACLs

Cisco security appliances provide basic traffic filtering capabilities with access rules and access control lists (ACLs), which control access in your network by preventing certain traffic from entering or exiting.

Note ASA documentation uses the terms "ACE" for access rule and "access list" for ACL.

This chapter describes how to add access rules to create ACLs, and it shows how to use ASDM to add them to your network configuration to control traffic.

- Information About Access Rules and ACLs, page 21-1
- Configuring Access Rules and ACLs, page 21-8

Information About Access Rules and ACLs

This section describes access rules and ACLs, and it includes the following topics:

- About Access Rules and ACLs, page 21-1
- About EtherType ACLs, page 21-3
- Allowing MPLS, page 21-4
- Using Access Rules, page 21-4
- Inbound and Outbound Access Rules and ACLs, page 21-4
- IP Addresses Used for Access Rules When You Use NAT, page 21-5
- Access Rules for Returning Traffic, page 21-7
- Allowing Broadcast and Multicast Traffic through the Transparent Firewall Using Access Rules, page 21-7

About Access Rules and ACLs

Access rules are used in a variety of features and can be configured for all routed and network protocols (IP, AppleTalk, and so on) to filter the packets of those protocols. Your access policy is made up of one or more access rules per interface. An access rule is a single entry in an ACL that specifies a permit or deny rule (to forward or drop the packet) and is applied to a protocol, to a source and destination IP address or network, and, optionally, to the source and destination ports.

You can configure the following types of access rules and ACLs:

- Standard Access Rules—Standard ACLs identify the destination IP addresses (not source addresses) of OSPF routes and can be used in a route map for OSPF redistribution. Standard ACLs cannot be applied to interfaces to control traffic. To configure standard access rules, see the "Configuring Standard ACLs" section on page 21-8.
- Extended Access Rules—An extended ACL is made up of one or more access rules in which you can specify the line number to insert the access rule, both the source and destination addresses, and, depending upon the access rule type, the protocol, the ports (for TCP or UDP), or the ICMP type (for ICMP). You can identify all of these parameters by creating an ACL, or you can use object groups and service groups for each parameter. (For more information about network objects and service groups, see the "Using Network Objects and Groups" section on page 20-1 and the "Configuring Service Groups" section on page 20-5. To configure extended access rules, see the "Configuring Extended ACLs" section on page 21-10.
- Webtype Access Rules—Webtype access rules are added to a configuration that supports filtering for clientless SSL VPN. To configure webtype access rules, see the "Configuring Webtype ACLs" section on page 21-14.
- EtherType Access Rules—For transparent mode only, Ethertype access rules are 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. See "Configuring EtherType ACLs" section on page 21-19 for more information.
- IPv6 Access Rules—ACL functionality in IPv6 is similar to typical ACLs in IPv4. Extended ACLs, Webtype ACLs, and EtherType ACLs allow you to define IPv6 access lists and set their deny and permit conditions using the IPv6 address of the desired interface. IPv6 does not support standard ACLs.

To access the security appliance interface for management access, you do not need an access rule allowing the host IP address. You only need to configure management access according to Chapter 18, "Configuring Management Access."

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

Note To allow any traffic to enter the security appliance, you must attach an inbound access rule to an interface; otherwise, the security appliance automatically drops all traffic that enters that interface.

About EtherType ACLs

This section includes the following topics:

- Implicit Deny in Ethertype Rules, page 21-3
- Supported EtherTypes, page 21-3
- Implicit Permit of IP and ARPs Only, page 21-3

Implicit Deny in Ethertype Rules

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

For EtherType rules, the implicit deny does not affect IPv4 traffic or ARPs; for example, if you allow EtherType 8037 (the EtherType for IPX), the implicit deny at the end of the list does not block any IP traffic that you previously allowed with an access rule (or implicitly allowed from a high security interface to a low security interface). However, if you explicitly deny all traffic with an EtherType rule, then IP and ARP traffic is denied.

Supported EtherTypes

An EtherType rule controls any EtherType identified by a 16-bit hexadecimal number.

EtherType rules support Ethernet V2 frames.

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

BPDUs, which are handled by the rule, are the only exception: they are SNAP-encapsulated, and the security appliance is designed to specifically handle BPDUs.

The security appliance receives trunk port (Cisco proprietary) BPDUs. Trunk BPDUs have VLAN information inside the payload, so the security appliance modifies the payload with the outgoing VLAN if you allow BPDUs.

Note If you use failover, you must allow BPDUs on both interfaces with an EtherType rule to avoid bridging loops.

Implicit Permit of IP and ARPs Only

IPv4 traffic is allowed through the transparent firewall automatically, without a rule, if it arrives from a higher security interface to a lower security interface. ARPs are allowed through the transparent firewall in both directions without a rule. ARP traffic can be controlled by ARP inspection.

However, to allow any traffic with EtherTypes other than IPv4 and ARP, you need to apply an EtherType access list, even from a high security to a low security interface.

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

Allowing MPLS

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

Using Access Rules

The following general information applies to using access rules:

- Same interface—You can apply access rules to each direction of an interface.
- Rule order—The order of rules is important. When the security appliance decides whether to forward or drop a packet, the security appliance tests the packet against each rule in the order in which the rules are listed. After a match is found, no more rules are checked. For example, if you create an access rule at the beginning of the list, and that rule explicitly permits all traffic for an interface, no further rules are ever checked.
- Disabling—You can disable a rule by making it inactive.
- Implicit deny—An ACL has an implicit deny at the end of the list, so unless you explicitly permit it, traffic cannot pass. For example, if you want to allow all users to access a network through the security appliance except for particular addresses, then you need to deny the particular addresses and then permit all others.

Inbound and Outbound Access Rules and ACLs

By default, all traffic from a higher-security interface to a lower-security interface is allowed. ACLs either let you allow traffic from lower-security interfaces or restrict traffic from higher-security interfaces.

The security appliance supports two types of ACLs:

- Inbound—Inbound ACLs apply to traffic as it enters an interface.
- Outbound—Outbound ACLs apply to traffic as it exits an interface.

Note The terms “inbound” and “outbound” refer to the application of an ACL on an interface, either to traffic entering the security appliance on an interface or traffic exiting the security appliance on an interface. These terms do not refer to the movement of traffic from a lower security interface to a higher security interface, commonly known as inbound, or from a higher to lower interface, commonly known as outbound.

An outbound ACL is useful, for example, if you want to allow only certain hosts on the inside networks to access a web server on the outside network. Rather than creating multiple inbound ACLs to restrict access, you can create a single outbound ACL that allows only the specified hosts. (See Figure 21-1.) The outbound ACL prevents any other hosts from reaching the outside network.

Figure 21-1 Outbound ACLs

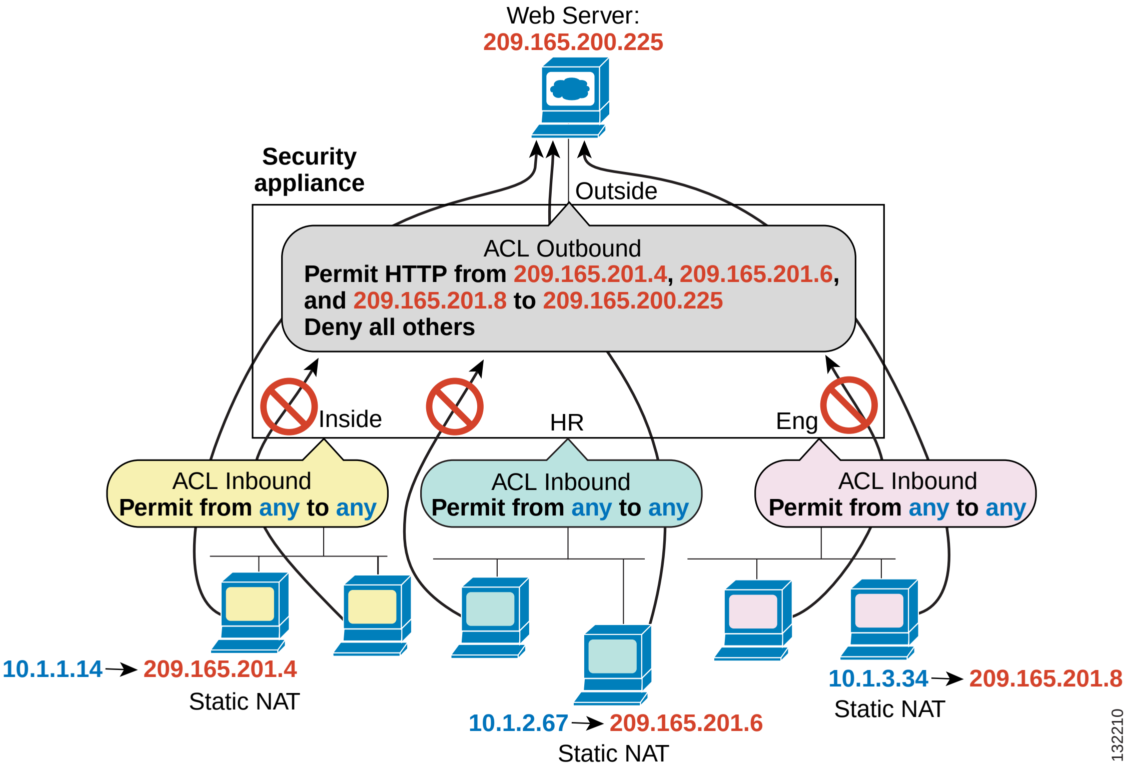

IP Addresses Used for Access Rules When You Use NAT

When you use NAT, the IP addresses you specify for an access rule depend on the interface to which the access rule is attached; you need to use addresses that are valid on the network that is connected to the interface. This guideline applies for both inbound and outbound access rules: the direction does not determine the address used, only the interface does.

For example, if you want to apply an access rule to the inbound direction of the inside interface, you configure the security appliance to perform NAT on the inside source addresses when they access outside addresses. Because the access rule is applied to the inside interface, the source addresses are the original untranslated addresses. Because the outside addresses are not translated, the destination address used in the access rule is the real address. (See Figure 21-2.)

Figure 21-2 IP Addresses in ACLs: NAT Used for Source Addresses

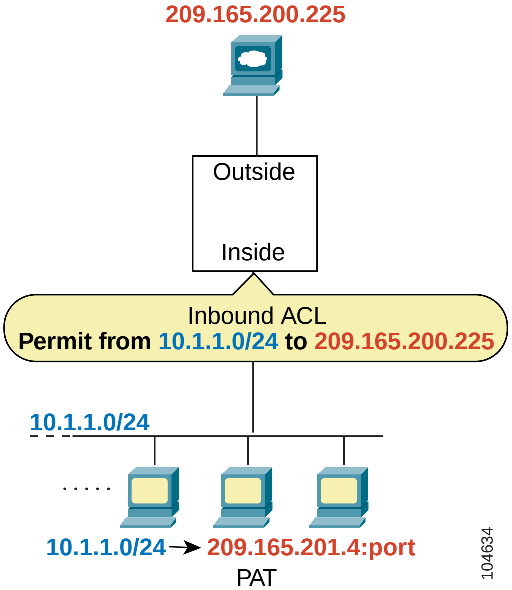
209.165.200.225
Outside
Inside
Inbound ACL
Permit from 10.1.1.0/24 to 209.165.200.225
10.1.1.0/24
.....
10.1.1.0/24 → 209.165.201.4:port
PAT
104634

If you want to allow an outside host to access an inside host, you can apply an inbound access rule on the outside interface. You need to specify the translated address of the inside host in the access rule because that address is the address that can be used on the outside network. (See Figure 21-3.)

Figure 21-3 IP Addresses in ACLs: NAT Used for Destination Addresses.

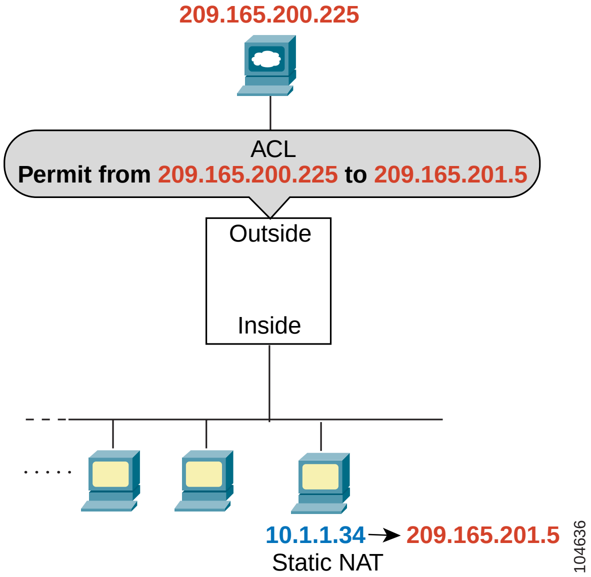

If you perform NAT on both interfaces, keep in mind the addresses that are visible to a given interface. In Figure 9-4, an outside server uses static NAT so that a translated address appears on the inside network.

Figure 21-4 IP Addressed in ACLs: NAT Used for Source and Destination Addresses

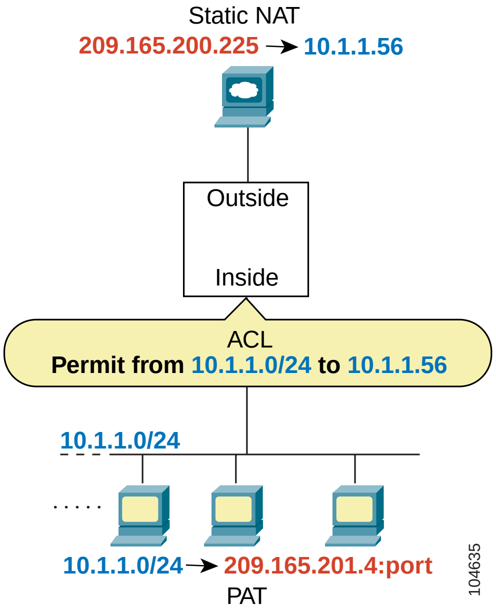

Access Rules for Returning Traffic

For TCP and UDP connections for both routed and transparent mode, you do not need an ACL to allow returning traffic because the security appliance allows all returning traffic for established, bidirectional connections. For connectionless protocols such as ICMP, however, the security appliance establishes unidirectional sessions, so you either need ACLs to allow ICMP in both directions (by applying access rules to the source and destination interfaces) or you need to enable the ICMP inspection engine. The ICMP inspection engine treats ICMP sessions as bidirectional connections. For more information about allowing ICMP in ACLs, see the "Adding an Access Rule to Manage a Service Group" section on page 21-20.)

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

In routed firewall mode, broadcast and multicast traffic is blocked even if you allow it in an access rule, including unsupported dynamic routing protocols and DHCP (unless you configure DHCP relay). Transparent firewall mode can allow any IP traffic through. This feature is especially useful in multiple context mode, which does not allow dynamic routing.

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

Table 21-1 lists common traffic types that you can allow through the transparent firewall.

Table 21-1 Transparent Firewall Special Traffic

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

Configuring Access Rules and ACLs

The Access Rules pane shows your entire network security policy expressed in rules.

When you choose the Access Rules option, this pane lets you define access rules to control the access of a specific host or network to another host/network, including the protocol or port that can be used.

For more information about access rules, see the "Information About Access Rules and ACLs" section on page 21-1.

This section includes the following topics:

- Configuring Standard ACLs, page 21-8
- Configuring Extended ACLs, page 21-10
- Configuring Webtype ACLs, page 21-14
- Configuring EtherType ACLs, page 21-19
- Adding an Access Rule to Manage a Service Group, page 21-20
- ASDM Field Definitions for Using Access Rules, page 21-23

Configuring Standard ACLs

Standard ACLs identify the destination IP addresses (not source addresses) of OSPF routes and can be used in a route map for OSPF redistribution. Standard ACLs cannot be applied to interfaces to control traffic. You must first create an ACL and then add an ACE to that ACL.

Note IPv6 does not support standard ACLs.

This section includes the following topics:

- Adding a Standard ACL, page 21-9
- Adding an ACE to a Standard ACL, page 21-9

Adding a Standard ACL

To add a standard ACL to your configuration and then add an ACE to the ACL, perform the following steps:

Step 1 Choose **Configuration** > **Firewall** > **Advanced** > **Standard ACL**.

Step 2 Click **Add**, and from the drop-down list, choose **Add ACL**.

Step 3 In the Add ACL dialog box, add a name or number (without spaces) to identify the ACL.

Step 4 Click **OK**.

The ACL name appears in the main pane.

You may add additional ACLs.

Step 5 Click **Apply** to save the ACLs to your configuration.

You can now add one or more ACEs to the newly created ACL.

To add an ACE, see the "Adding an ACE to a Standard ACL" section on page 21-9.

Adding an ACE to a Standard ACL

Before you can add an ACE to a configuration, you must first add an ACL. For information about adding a standard ACL, see the "Adding a Standard ACL" section on page 21-9. For information about editing ACEs, see the "Editing an ACE in a Standard ACL" section on page 21-10.

To add an ACE to an ACL that exists in your configuration, perform the following steps:

Step 1 Choose **Configuration** > **Firewall** > **Advanced** > **Standard ACL**.

Step 2 In the main pane, select the ACL for which you want to add an ACE.

Step 3 Click **Add**, and choose **Add ACE** from the drop-down list.

The Add ACE dialog box appears.

Step 4 (Optional) To specify the placement of the new ACE, select an existing ACE, and click **Insert...** to add the ACE before the selected ACE, or click **Insert After...** to add the ACE after the selected ACE.

Step 5 Click one of the following radio buttons to choose an action:

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

Step 6 In the Address field, enter the IP address of the destination to which you want to permit or deny access.

You can also browse for the address of a network object by clicking the ellipsis at the end of the Address field.

Step 7 (Optional) In the Description field, enter a description that makes the ACE easier to understand.

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

Step 8 Click **OK**.

The newly created ACE appears under the ACL.

Step 9 Click **Apply** to save the ACE to your configuration.

Editing an ACE in a Standard ACL

To edit an ACE in a standard ACL, perform the following steps:

Step 1 Choose **Configuration** > **Firewall** > **Advanced** > **Standard ACL**.

Step 2 In the main pane, select the ACE that you want to edit.

Step 3 Click **Edit**,

The Edit ACE dialog box appears.

Step 4 Enter the desired changes.

Step 5 Click **OK**.

Configuring Extended ACLs

Extended ACLs have the ability to filter packets based on source and destination IP addresses. An extended ACL is made up of one or more access rules in which you can specify the line number to insert the access rule, both the source and destination addresses, and, depending upon the access rule type, the protocol, the ports (for TCP or UDP), or the ICMP type (for ICMP). You can identify all of these parameters by creating an ACL, or you can use object groups and service groups for each parameter. (For more information about network objects and service groups, see the "Using Network Objects and Groups" section on page 20-1 and the "Configuring Service Groups" section on page 20-5.

This section includes the following topics:

- Configuring Extended ACLs for Management Traffic, page 21-10
- Configuring Extended Access Rules for Network Traffic, page 21-12
- Editing Extended ACLs for Management Traffic, page 21-12
- Editing Extended Access Rules for Network Traffic, page 21-13
- Deleting an Extended ACEs, page 21-14

Configuring Extended ACLs for Management Traffic

You can configure an interface ACL that supports access control for to-the-box management traffic from a specific peer (or set of peers) to the security appliance. One scenario in which this type of ACL would be useful is when you want to block IKE Denial of Service attacks. (To edit extended ACLs for management traffic, see the "Editing Extended ACLs for Management Traffic" section on page 21-12.)

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

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

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

- **Add Management Access Rule**
- **Add IPv6 Management Access Rule**

The appropriate access rule dialog box appears.

Step 3 From the Interface drop-down list, choose the interface on which to apply the rule.

The management interface is for management only and cannot be used to configure an access rule.

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

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

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

Note IPv6 must be enabled on at least one interface before you can configure an extended ACL with an IPv6 address. For more information about enabling IPv6 on an interface, see Chapter 9, "Configuring Interfaces."

Step 6 Select the service type.

For more information about service types, see the "Adding an Access Rule to Manage a Service Group" section on page 21-20.

Step 7 (Optional) In the Description field, add a text description about the ACL.

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

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

For more information about logging options, see the "Log Options" section on page 21-29.

Step 9 (Optional) To add a source service (TCP, UDP, and TCP-UDP only) and a time range to your access rule that specifies when traffic can be allowed or denied, click **More Options** to expand the list.

a. In the Source Service field, click the browse button.

b. In the Browse Source Services window, choose the service to apply to the access rule, and click **Source Service**.

The selected service appears in the Source Service field.

c. Click **OK**.

Step 10 (Optional) To add a time range, perform the following steps under More Options.

a. To the right of the Time Range field, click the browse button.

The Browse Time Range dialog box appears.

b. Click **Add**.

The Add Time Range dialog box appears.

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

d. Choose the Start Time and the End Time.

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

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

Step 11 Click **Apply** to save the ACL to your configuration.

Note After you add access rules, you can click the following radio buttons to filter which access rules appear in the main window: IPv4 and IPv6, IPv4 Only, or IPv6 Only.

Editing Extended ACLs for Management Traffic

To edit an extended ACL for management traffic, perform the following steps:

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

Step 2 Click **Edit**.

The Edit Management Access Rule dialog box appears.

Step 3 Enter the desired changes.

For information about fields in the ACL, see the "Configuring Extended ACLs for Management Traffic" section on page 21-10 or the "ASDM Field Definitions for Using Access Rules" section on page 21-23.

Step 4 Click **OK**.

Step 5 Click **Apply** to save the changes to your configuration.

Configuring Extended Access Rules for Network Traffic

Extended access rules can also control access through-the box, filtering from hosts through the security appliance and to the outside network. This section shows how to add an extended access rule. To edit an extended ACL for network traffic, see the "Editing Extended Access Rules for Network Traffic" section on page 21-13.

To configure an extended access rule for network traffic, perform the following steps:

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

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

- **Add Access Rule**
- **Add IPv6 Access Rule**

The appropriate access rule dialog box appears.

Step 3 From the Interface drop-down list, choose the interface on which to apply the rule.

The management interface is for management only and cannot be used to configure an access rule.

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

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

Step 5 In the Source field, enter an IP address that specifies the network object group, interface IP, or any address from which traffic is permitted or denied to the specified destination.

For more information about enabling IPv6 on an interface, see Chapter 9, "Configuring Interfaces."

Step 6 In the Destination field, enter an IP address that specifies the network object group, interface IP, or any address to which traffic is permitted or denied from the source specified in the Source field.

Step 7 Select the service type.

For more information about service types, see the "Adding an Access Rule to Manage a Service Group" section on page 21-20.

Step 8 (Optional) To add a time range to your access rule that specifies when traffic can be allowed or denied, click **More Options** to expand the list.

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

The Browse Time Range dialog box appears.

b. Click **Add**.

The Add Time Range dialog box appears.

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

d. Choose the Start Time and the End Time.

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

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

Step 9 (Optional) In the Description field, add a text description about the ACL.

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

Step 10 (Optional) Logging is enabled by default. You can disable logging by unchecking the check box, or you ca change the logging level from the drop-down list. The default logging level is Informational.

For more information about logging options, see the "Log Options" section on page 21-29.

Step 11 Click **OK**. The ACL appears with the newly configured access rules.

Step 12 Click **Apply** to save the ACL to your configuration.

Note After you add access rules, you can click the following radio buttons to filter which access rules appear in the main pane: IPv4 and IPv6, IPv4 Only, or IPv6 Only.

Editing Extended Access Rules for Network Traffic

To edit an extended access rule for network traffic, perform the following steps:

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

Step 2 Choose the Access Rule Type to edit by clicking one of the following radio button:

- **IPv4 and IPv6**—Shows access rules that have both IPv4 and IPv6 type addresses.

- **IPv4 Only**—Shows access rules that have IPv4 type addresses only.
- **IPv6 Only**—Shows access rules that have IPv6 type addresses only.

The main Access Rule pane displays the available interfaces for the chosen rule type.

Step 3 Select the ACE to edit.

Step 4 Click **Edit**.

The Edit Access Rule dialog box appears.

Step 5 Enter changes to the current configuration.

Step 6 Click **OK**.

The main Access Rules pane displays the updated access rules.

Step 7 Click **Apply** to save the changes to your configuration.

Deleting an Extended ACEs

To delete an extended ACE, perform the following steps:

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

Step 2 Choose the Access Rule Type to edit by clicking one of the following radio buttons:

- **IPv4 and IPv6**—Shows access rules that have both IPv4 and IPv6 type addresses.
- **IPv4 Only**—Shows access rules that have IPv4 type addresses only.
- **IPv6 Only**—Shows access rules that have IPv6 type addresses only.

Step 3 Select the existing ACE to delete.

Step 4 Click **Delete**.

The main Access Rules pane displays without the selected ACE.

Step 5 Click **Apply** to save the configuration.

Configuring Webtype ACLs

Webtype ACLs are added to a configuration that supports filtering for clientless SSL VPN. The ACL permits or denies access to specific networks, subnets, hosts, and web servers. Each ACE specifies one rule that serves the function of the ACL. If you do not configure any filters, all connections are permitted.

This section includes the following topics:

- Adding a Webtype ACL and ACE, page 21-15
- Examples, page 21-16
- Editing Webtype ACLs and ACEs, page 21-18

Adding a Webtype ACL and ACE

You must first create the webtype ACL and then add an ACE to the ACL.

Note Smart tunnel ACEs filter only on a per-server basis. Thus, you cannot create smart tunnel ACEs to do the following:

- Permit or deny access to directories.
- Permit or deny access to specific smart tunnel-enabled applications.

To configure a webtype access rule, perform the following steps:

Step 1 Choose **Configuration** > **Remote Access VPN** > **Clientless SSL VPN Access** > **Advanced** > **Web ACLs**.

Step 2 Click **Add**, and choose one of the following ACL types to add:

- **Add ACL**
- **Add IPv6 ACL**

The Add ACL dialog box appears.

Step 3 Enter a name for the ACL (with no spaces), and click **OK**.

Step 4 To add an entry to the list that you just created, click **Add**, and choose **Add ACE** from the drop-down list.

Step 5 In the Action field, click the radio button next to the desired action:

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

Note The end of every ACL has an implicit deny rule.

Step 6 In the Filter field, you can either filter on a URL or filter on an address and Service.

a. To filter on a URL, choose the URL prefix from the drop-down list, and enter the URL.

Wild card characters can be used in the URL field:

– An asterisk * matches none or any number of characters.

– A question mark (?) matches any one character exactly.

– Square brackets [] are range operators, matching any character in the range. For example, to match both http://www.cisco.com:80/ and http://www.cisco.com:81/, enter the following:

http://www.cisco.com:8[01]/

b. To filter on an address and service, click the **Filter address and service** radio button, and enter the appropriate values.

You can also browse for the address and service by clicking the browse buttons at the end of the fields.

Step 7 (Optional) Logging is enabled by default. You can disable logging by unchecking the check box, or you can change the logging level from the drop-down list. The default logging level is Informational.

For more information about logging options, see the "Log Options" section on page 21-29.

Step 8 (Optional) If you changed the logging level from the default setting, you can specify the logging interval by clicking **More Options** to expand the list.

Valid values are from 1 through 6000 seconds. The default is 300 seconds.

Step 9 (Optional) To add a time range to your access rule that specifies when traffic can be allowed or denied, click **More Options** to expand the list.

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

The Browse Time Range dialog box appears.

b. Click **Add**.

The Add Time Range dialog box appears.

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

d. Enter the Start Time and the End Time.

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

Step 10 Click **OK** to apply the optional time range specifications.

Step 11 Click **Apply** to save the configuration.

Note After you add access rules, you can click the following radio buttons to filter which access rules appear in the main pane: IPv4 and IPv6, IPv4 Only, or IPv6 Only.

Examples

See the following example ACEs:

- Example Smart Tunnel ACEs
- Example Port Forwarding TCP-based ACLs
- Example Plug-in ACEs
- Example HTTP and HTTPS ACEs
- Example CIFS Shares ACEs

Table 21-2 Example Smart Tunnel ACEs

		Filter on URL Values	
Function	**Action**	**drop-down**	**:// text box**
Permit a basic URL.	Permit	any	http://www.example.com
Permit smart tunnel access to a basic URL. **Note**: To configure smart tunnel access, add both ACE types, then choose Clientless SSL VPN Access > Portal > Bookmarks, specify the URL in a bookmark entry, and check Enable Smart Tunnel in the Add Bookmark dialog box.	Permit	any	http://www.example.com
	Permit	any	smart-tunnel://www.example.com

Table 21-2 Example Smart Tunnel ACEs

Function	Action	Filter on URL Values	
		drop-down	:// text box
Permit smart tunnel access to a basic URL, but deny smart tunnel access to another.	Permit	any	http://www.example.com
	Permit	any	smart-tunnel://www.example.com
	Deny	any	smart-tunnel://images.example.com
Permit smart tunnel access to all URLs ending with ".example.com"	Permit	any	http://*.example.com
	Permit	any	smart-tunnel://*.example.com

Table 21-3 Example Port Forwarding TCP-based ACLs

Function	Action	Filter on Address and Service Values	
		Address	Service
Permit port forwarding.	Permit	192.168.20.92	2224
Block ports 80 to 90.	Deny	192.168.20.92	80-90
Block all ports.	Deny	192.168.20.92	tcp
Implicitly deny all traffic except that destined to the address specified.	Permit	10.86.192.1	tcp
Specify the default TCP port (3389).	Permit	10.86.192.193	3389

Note Protocols such as RDP, SSH, and VNC are available only if the respective plug-in is imported on the security appliance.

Table 21-4 Example Plug-in ACEs

Function	Action	Filter on URL Values		Filter on Address and Service Values	
		drop-down	:// text box	Address	Service
Permit SSH plug-in access to the specified IP address.	Permit	ssh	192.168.20.92		
Permit RDP plug-in access to a range of IP addresses.	Permit	rdp	192.168.20.[1-112]		
Use a wildcard to specify a range of IP addresses to permit RDP plug-in access.	Permit	rdp	192.168.20.*		
Provide access to specific URLs.	Permit	http	10.80.192.1/engineering/*		
	Permit	http	10.86.192.1/marketing/*		
	Permit			10.86.192.193	3389
Provide TCP access.	Permit			10.86.192.193	3389

Table 21-5 Example HTTP and HTTPS ACEs

Function	Action	Filter on URL Values	
		drop-down	**:// text box**
Provide access to specific URLs.	Permit	http	10.80.192.1/engineering/*
	Permit	http	10.86.192.1/marketing/*
Provide access to any HTTP URL.	Permit	http	*/*
Provide access to any HTTPS URL.	Permit	https	*/*

Table 21-6 Example CIFS Shares ACEs

Function	Action	Filter on URL Values	
		drop-down	**:// text box**
Deny access to the shares/Marketing_Reports folder. If this entry were the only one in the ACL for the IP address specified, it would implicitly deny access to the root and all peer and sub-folders of shares/Sales_Reports.	Deny	cifs	172.16.10.39/shares/Marketing_Reports
Permit access to all remaining folders of the specified IP address.	Permit	cifs	172.16.10.40/shares*

Editing Webtype ACLs and ACEs

To edit a webtype ACL or ACE, perform the following steps:

Step 1 Choose **Configuration** > **Remote Access VPN** > **Clientless SSL VPN Access** > **Advanced** > **Web ACLs**.

Step 2 Choose the Access Rule Type to edit by clicking one of the following radio buttons:

- **IPv4 and IPv6**—Shows access rules that have both IPv4 and IPv6 type addresses.
- **IPv4 Only**—Shows access rules that have IPv4 type addresses only.
- **IPv6 Only**—Shows access rules that have IPv6 type addresses only.

The main Access Rule pane displays the available interfaces for the chosen rule type.

Step 3 Select the ACE to edit, and make any changes to the values.

For more information about specific values, see the "Adding a Webtype ACL and ACE" section on page 21-15 or the "ASDM Field Definitions for Using Access Rules" section on page 21-23. The "Examples" section on page 21-16 shows example uses of ACEs.

Step 4 Click **OK**.

Step 5 Click **Apply** to save the changes to your configuration.

Deleting Webtype ACEs

To delete a webtype ACE, perform the following steps:

Step 1 Choose **Configuration** > **Remote Access VPN** > **Clientless SSL VPN Access** > **Advanced** > **Web ACLs**.

Step 2 Choose the Access Rule Type of the rule you want to delete by clicking the radio button:

- IPv4 and IPv6—Shows access rules that have both IPv4 and IPv6 type addresses.
- IPv4 Only—Shows access rules that have IPv4 type addresses only.
- IPv6 Only—Shows access rules that have IPv6 type addresses only.

The main Access Rule pane displays the available interfaces for the chosen rule type.

Step 3 Select the ACE that you want to delete.

Note If you select a specific ACE, only that ACE is deleted. If you select an ACL, that ACL and all of the ACEs under it are deleted.

Step 4 Click **Delete**.

Step 5 The selected items are removed from the viewing pane.

Note If you deleted a rule in error and want to restore it to your configuration, click **Reset** before you click Apply. The deleted rules reappear in the viewing pane.

Step 6 Click **Apply** to save the change to the configuration.

Configuring EtherType ACLs

Ethertype access rules are 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.

For information about EtherType access rules and ACLs, see the "About EtherType ACLs" section on page 21-3.

This section includes the following topics:

- Adding EtherType ACLs, page 21-19
- Editing EtherType Access Rules, page 21-20

Adding EtherType ACLs

To add an EtherType ACL, perform the following steps:

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

Step 2 Click **Add**.

The Add EtherType rules window appears.

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

Step 4 From the Interface drop-down list, choose the interface on which to apply the rule.

The management interface is for management only and cannot be used to configure an access rule.

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

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

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

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

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

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

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

Step 9 Click **OK**.

Editing EtherType Access Rules

You can edit the contents of EtherType access rules columns to change a rule or to sort its order in the configuration.

To edit an EtherType ACL, perform the following steps:

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

Step 2 Click **Edit**.

The Edit EtherType Rule dialog box appears.

Step 3 Enter the desired changes, and click **OK**.

For specific information about the access rule fields, see the "Adding EtherType ACLs" section on page 21-19.

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

Adding an Access Rule to Manage a Service Group

The Service Groups pane lets you associate multiple services into a named group. You can specify any type of protocol and service in one group, or 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 configure access rules, the ASDM window includes a Services pane at the right that shows available service groups and other global objects. You can add, edit, or delete objects directly in the Services pane. For more information about managing service groups, see the “Manage Service Groups” section on page 21-27.

To configure an ACL with a service group, perform the following steps:

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

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

- **Add Access Rule**
- **Add IPv6 Access Rule**

The appropriate access rule dialog box appears.

Step 3 From the Interface drop-down list, choose the interface on which to apply the rule.

The management interface is for management only and cannot be used to configure an access rule.

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

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

Step 5 In the Source field, enter an IP address that specifies the network object group, interface IP, or any address from which traffic is permitted or denied to the specified destination.

For more information about enabling IPv6 on an interface, see Chapter 9, “Configuring Interfaces.”

Step 6 In the Destination field, enter an IP address that specifies the network object group, interface IP, or any address to which traffic is permitted or denied from the source specified in the Source field.

Step 7 In the Service field, click the browse button.

The Browse Service dialog box appears.

Step 8 From the list Browse Service list, select the service group for which you want to apply the access rule. You can choose either TCP ports, UDP ports, TCP-UDP ports, ICMP types, or IP protocols.

Step 9 Click **OK**.

Step 10 (Optional) In the Description pane, add a description of the access rule (up to 200 characters in length).

Step 11 (Optional) Logging is enabled by default. You can disable logging by unchecking the check box, or you can change the logging level from the drop-down list. The default logging level is Informational.

For more information about logging options, see the “Log Options” section on page 21-29.

Step 12 (Optional) To add a source service (TCP, UDP, and TCP-UDP only) and a time range to your access rule that specifies when traffic can be allowed or denied, click **More Options** to expand the list.

a. In the Source Service field, click the browse button to browse for an existing service.

b. In the Browse Source Services dialog box, choose the service to apply to the access rule, and click the Source Service button.

The selected service appears in the Source Service field.

c. Click **OK**.

Step 13 (Optional) To add a time range, perform the following steps under More Options.

a. To the right of the Time Range field, click the browse button.

The Browse Time Range dialog box appears.

b. Click **Add**.

The Add Time Range dialog box appears.

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

d. Enter the Start Time and the End Time.

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

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

Step 14 Click **Add**.

Step 15 From the Add drop-down list, choose **Service Group**.

The Add Service Group dialog box appears.

Step 16 In the Group Name field, enter a meaningful name for the service group.

Step 17 Click the **Existing Service/Service Group** radio button.

Step 18 Click **OK**.

The Browse Service pane displays the newly configured service group.

Step 19 Click **OK**.

The Service field appears with the newly configured service group.

Step 20 Click **OK**.

Step 21 Click **Apply** to save the configuration.

Note After you add access rules, you can click one of the following radio buttons to filter which access rules appear in the main pane: **IPv4 and IPv6**, **IPv4 Only**, or **IPv6 Only**.

ASDM Field Definitions for Using Access Rules

This section includes the following topics:

Access Rule Field Definitions

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 that 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, either Interface, Source, Destination, Source or Destination, Destination Service, or Rule Query. A rule query is a collection of multiple criteria that you can save and use repeatedly.
  - Condition drop-down list—For criteria Source, Destination, Source or Destination, and Destination Service, choose the condition, either “is” or “includes.”
  - Filter field—The Interface type field becomes a drop-down list so that 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 an address manually or browse for one by clicking the browse button (...) and launching the Browse Address dialog box. The Destination Service type accepts a TCP, UDP, TCP-UDP, ICMP, or IP protocol type. You can enter a type manually or browse for one by clicking the browse 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 that you can manage named rule queries.
- 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.
- Export—Exports to a file in either comma separated value or HTML format.
- Clear Hits—Clears the counted hits for the selected access rule. Logging must be enabled for this field to be active.
- Show Log—Shows the syslogs generated by the selected access rule in the Real-Time Log Viewer.
- Packet Trace—Provides detailed information about packet processing with the adaptive security appliance, as well as information for packet sniffing and network fault isolation.
- IPv4 Only—Shows access rules that have IPv4 type addresses only.
- IPv6 Only—Shows access rules that have IPv6 type addresses only.
- IPv4 and IPv6—Shows access rules that have both IPv4 and IPv6 type addresses.

The following description summarizes the columns in the Access Rules table. You can edit the contents of these columns by double-clicking 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 address 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 configuration 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 address 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 configuration 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 it 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 rule. 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.
- Hits—Shows the number of hits for the rule. This column is dynamically updated depending upon the frequency set in the Preferences dialog box. Hit counts are applicable for explicit rules only. No hit count will be displayed for implicit rules in the Access Rules table.
- Logging—Shows the logging level and the interval in seconds between log messages (if you enable logging for the access rule).
- 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—Enables you to 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—Enables you to add, edit, delete, or find services.
- Time Ranges—Enables you to add, edit, or delete time ranges.

Modes

The following table shows the modes in which ACLs are available:

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

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

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—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—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 browse (...) 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 browse (...) 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

The Add/Edit Rule dialog box lets you create a new rule or modify an existing rule.

For more information about access rules, see the "Information About Access Rules and ACLs" section on page 21-1.

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.
- 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 field.

...—Lets you select, add, edit, delete, or find an existing IP address object, IP name, network object group, or all.

- 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.

  ...—Lets you select, add, edit, delete, or find an existing IP address object, IP name, network object group, or all.

- Service—Choose this option to specify a port number, a range of ports, or a well-known service name or group from a list of services.

  ...—Lets you select, add, edit, delete, or find an existing service from a preconfigured list.

- Description—(Optional) Enter a description of the access rule.
- Enable Logging—Enables logging for the access rule.
  - Logging Level—Specifies default, emergencies, alerts, critical, errors, warnings, notifications, informational, or debugging.
- More Options—Shows additional configuration options for the rule.
  - Enable Rule—Enables or disables the rule.
  - Traffic Direction—Determines which direction of traffic the rule is applied. Options are either incoming or outgoing.
  - Source Service—Specifies a source protocol and service (TCP or UDP service only).

    ...—Lets you select, add, edit, delete or find a source service from a preconfigured list.
  - Logging Interval—Specifies the interval for logging in seconds if logging is configured.
  - Time Range—Specifies a time range defined for this rule from the drop-down list.

    ...—Lets you select, add, edit, delete or find a time range from a preconfigured list.

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

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—Choose this option to add TCP services or port numbers to an object group.
- UDP—Choose this option to add UDP services or port numbers to an object group.
- TCP-UDP—Choose 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 in this list, select the desired 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

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 choose 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

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

When you enable logging, if a packet matches the access rule, the 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 access rule 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.

For more information about access rules, see the "Information About Access Rules and ACLs" section on page 21-1.

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 rule. 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 rule, the access rule provided by a RADIUS server replaces the access rule configured on that interface. If the per user override feature is disabled, the access rule provided by the RADIUS server is combined with the access rule configured on that interface. If the inbound access rule is not configured for the interface, per user override cannot be configured.

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

The Log Options dialog box lets you set logging options for each access rule. See the "Advanced Access Rule Configuration" section on page 21-29 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 dialog box.

Fields

- Use default logging behavior—Uses the older access rule 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 rule logging mechanism: the security appliance logs system log message number 106100 when a packet matches the access rule (either permit or deny).

  If a packet matches the access rule, 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 access rule 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 access rule. The default is 300 seconds.
- Disable logging for the rule—Disables all logging for 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
•	•	•	•	—

CHAPTER 22

Configuring NAT

This chapter describes Network Address Translation, and includes the following sections:

- NAT Overview, page 22-1
- Configuring NAT Control, page 22-15
- Using Dynamic NAT, page 22-16
- Using Static NAT, page 22-26
- Using NAT Exemption, page 22-30

NAT Overview

This section describes how NAT works on the security appliance, and includes the following topics:

- Introduction to NAT, page 22-1
- NAT Control, page 22-4
- NAT Types, page 22-6
- Policy NAT, page 22-10
- NAT and Same Security Level Interfaces, page 22-12
- Order of NAT Rules Used to Match Real Addresses, page 22-13
- Mapped Address Guidelines, page 22-13
- DNS and NAT, page 22-13

Introduction to NAT

Address translation substitutes the real address in a packet with a mapped address that is routable on the destination network. NAT is composed of two steps: the process by which a real address is translated into a mapped address, and the process to undo translation for returning traffic.

The security appliance translates an address when a NAT rule matches the traffic. If no NAT rule matches, processing for the packet continues. The exception is when you enable NAT control. NAT control requires that packets traversing from a higher security interface (inside) to a lower security interface (outside) match a NAT rule, or processing for the packet stops. See the "Default Settings" section on page 9-7 for more information about security levels. See the "NAT Control" section on page 22-4 for more information about NAT control.

Note In this document, all types of translation are referred to as NAT. When describing NAT, the terms *inside* and *outside* represent the security relationship between any two interfaces. The higher security level is inside and the lower security level is outside. For example, interface 1 is at 60 and interface 2 is at 50; therefore, interface 1 is "inside" and interface 2 is "outside."

Some of the benefits of NAT are as follows:

- You can use private addresses on your inside networks. Private addresses are not routable on the Internet.
- NAT hides the real addresses from other networks, so attackers cannot learn the real address of a host.
- You can resolve IP routing problems such as overlapping addresses.

See Table 25-1 on page 25-3 for information about protocols that do not support NAT.

NAT in Routed Mode

Figure 22-1 shows a typical NAT example in routed mode, with a private network on the inside. When the inside host at 10.1.1.27 sends a packet to a web server, the real source address, 10.1.1.27, of the packet is changed to a mapped address, 209.165.201.10. When the server responds, it sends the response to the mapped address, 209.165.201.10, and the security appliance receives the packet. The security appliance then changes the translation of the mapped address, 209.165.201.10 back to the real address, 10.1.1.1.27 before sending it to the host.

Figure 22-1 NAT Example: Routed Mode

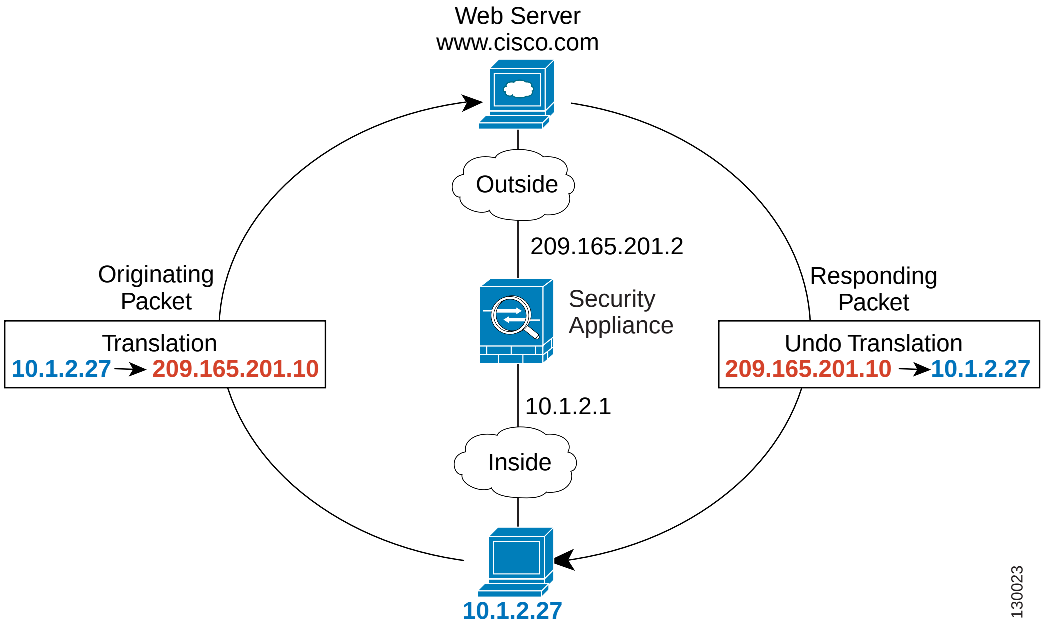

NAT in Transparent Mode

Using NAT in transparent mode eliminates the need for the upstream or downstream routers to perform NAT for their networks. For example, a transparent firewall security appliance is useful between two VRFs so you can establish BGP neighbor relations between the VRFs and the global table. However, NAT per VRF might not be supported. In this case, using NAT in transparent mode is essential.

NAT in transparent mode has the following requirements and limitations:

- When the mapped addresses are not on the same network as the transparent firewall, then on the upstream router, you need to add a static route for the mapped addresses that points to the downstream router (through the security appliance).
- If the real destination address is not directly-connected to the security appliance, then you also need to add a static route on the security appliance for the real destination address that points to the downstream router. Without NAT, traffic from the upstream router to the downstream router does not need any routes on the security appliance because it uses the MAC address table. NAT, however, causes the security appliance to use a route lookup instead of a MAC address lookup, so it needs a static route to the downstream router.
- The **alias** command is not supported.
- Because the transparent firewall does not have any interface IP addresses, you cannot use interface PAT.
- ARP inspection is not supported. Moreover, if for some reason a host on one side of the firewall sends an ARP request to a host on the other side of the firewall, and the initiating host real address is mapped to a different address on the same subnet, then the real address remains visible in the ARP request.

Figure 22-2 shows a typical NAT scenario in transparent mode, with the same network on the inside and outside interfaces. The transparent firewall in this scenario is performing the NAT service so that the upstream router does not have to perform NAT. When the inside host at 10.1.1.27 sends a packet to a web server, the real source address of the packet, 10.1.1.27, is changed to a mapped address, 209.165.201.10. When the server responds, it sends the response to the mapped address, 209.165.201.10, and the security appliance receives the packet because the upstream router includes this mapped network in a static route directed through the security appliance. The security appliance then undoes the translation of the mapped address, 209.165.201.10 back to the real address, 10.1.1.1.27. Because the real address is directly-connected, the security appliance sends it directly to the host.

Figure 22-2 NAT Example: Transparent Mode

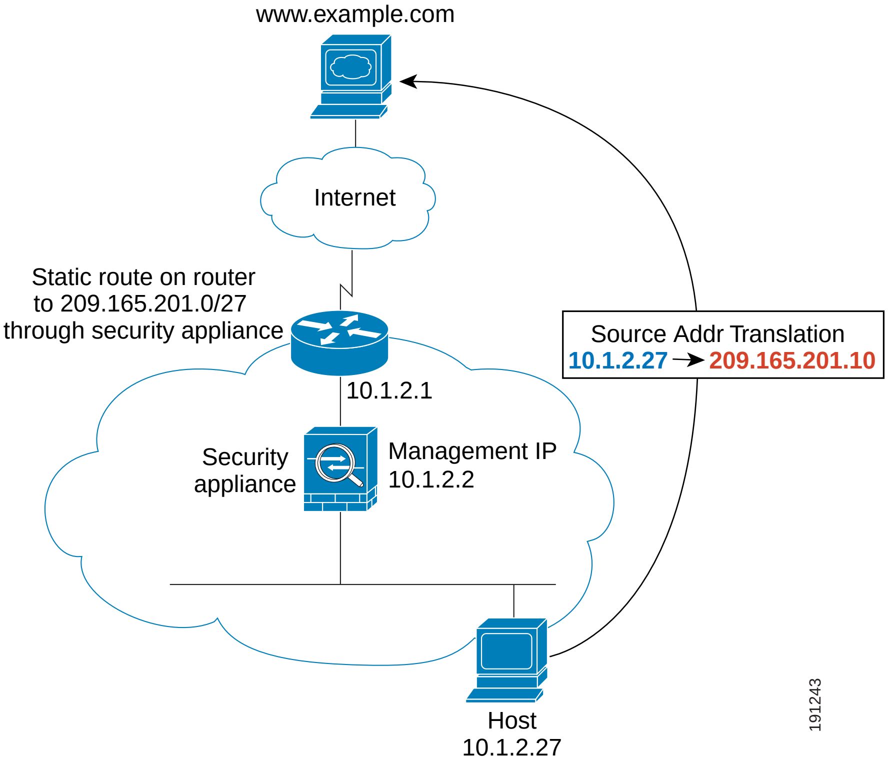

NAT Control

NAT control requires that packets traversing from an inside interface to an outside interface match a NAT rule; for any host on the inside network to access a host on the outside network, you must configure NAT to translate the inside host address, as shown in Figure 22-3.

Figure 22-3 NAT Control and Outbound Traffic

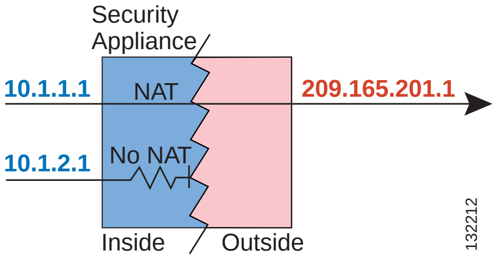

Interfaces at the same security level are not required to use NAT to communicate. However, if you configure dynamic NAT or PAT on a same security interface, then all traffic from the interface to a same security interface or an outside interface must match a NAT rule, as shown in Figure 22-4.

Figure 22-4 NAT Control and Same Security Traffic

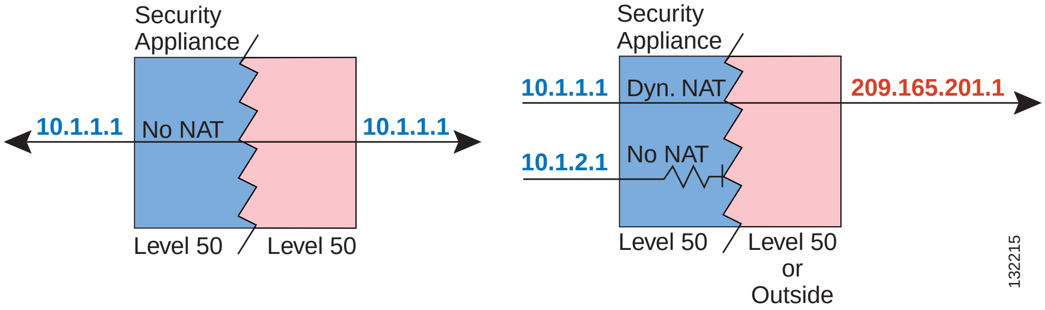

Similarly, if you enable outside dynamic NAT or PAT, then all outside traffic must match a NAT rule when it accesses an inside interface (see Figure 22-5).

Figure 22-5 NAT Control and Inbound Traffic

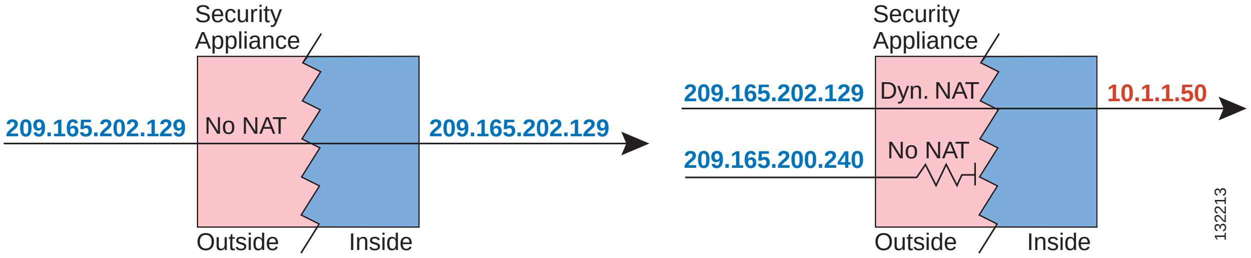

Static NAT does not cause these restrictions.

By default, NAT control is disabled; therefore, you do not need to perform NAT on any networks unless you want to do so. If you upgraded from an earlier version of software, however, NAT control might be enabled on your system. Even with NAT control disabled, you need to perform NAT on any addresses for which you configure dynamic NAT. See the "Dynamic NAT Implementation" section on page 22-16 for more information about how dynamic NAT is applied.

If you want the added security of NAT control but do not want to translate inside addresses in some cases, you can apply a NAT exemption or identity NAT rule on those addresses. (See the "Using NAT Exemption" section on page 22-30 for more information).

To configure NAT control, see the "Configuring NAT Control" section on page 22-15.

Note In multiple context mode, the packet classifier might rely on the NAT configuration to assign packets to contexts if you do not enable unique MAC addresses for shared interfaces. See the "How the Security Appliance Classifies Packets" section on page 11-2 for more information about the relationship between the classifier and NAT.

NAT Types

This section describes the available NAT types, and includes the following topics:

- Dynamic NAT, page 22-6
- PAT, page 22-8
- Static NAT, page 22-8
- Static PAT, page 22-9
- Bypassing NAT When NAT Control is Enabled, page 22-10

You can implement address translation as dynamic NAT, Port Address Translation, static NAT, static PAT, or as a mix of these types. You can also configure rules to bypass NAT; for example, to enable NAT control when you do not want to perform NAT.

Dynamic NAT

Dynamic NAT translates a group of real addresses to a pool of mapped addresses that are routable on the destination network. The mapped pool may include fewer addresses than the real group. When a host you want to translate accesses the destination network, the security appliance assigns the host an IP address from the mapped pool. The translation is added only when the real host initiates the connection. The translation is in place only for the duration of the connection, and a given user does not keep the same IP address after the translation times out. Users on the destination network, therefore, cannot initiate a reliable connection to a host that uses dynamic NAT, although the connection is allowed by an access list, and the security appliance rejects any attempt to connect to a real host address directly. See the "Static NAT" or "Static PAT" section for information on how to obtain reliable access to hosts.

Note In some cases, a translation is added for a connection, although the session is denied by the security appliance. This condition occurs with an outbound access list, a management-only interface, or a backup interface in which the translation times out normally.

Figure 22-6 shows a remote host attempting to connect to the real address. The connection is denied, because the security appliance only allows returning connections to the mapped address.

Figure 22-6 Remote Host Attempts to Connect to the Real Address

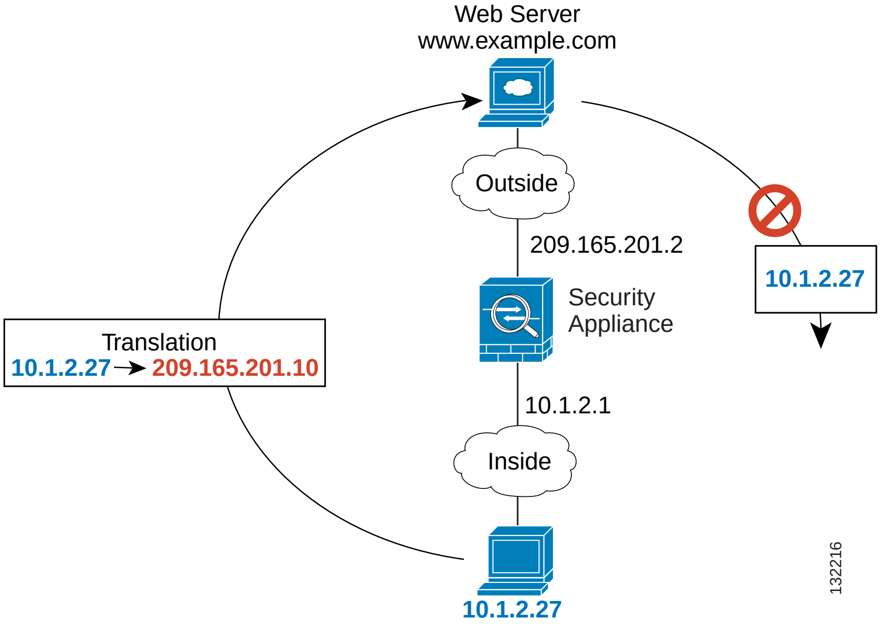

Figure 22-7 shows a remote host attempting to initiate a connection to a mapped address. This address is not currently in the translation table; therefore, the security appliance drops the packet.

Figure 22-7 Remote Host Attempts to Initiate a Connection to a Mapped Address

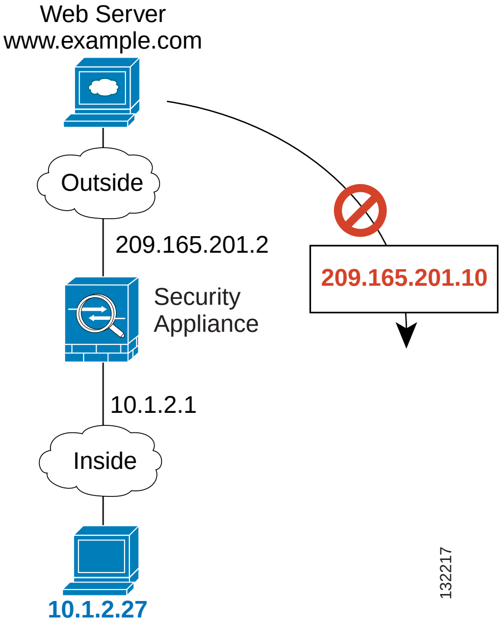

Note For the duration of the translation, a remote host can initiate a connection to the translated host if an access list allows it. Because the address is unpredictable, a connection to the host is unlikely. Nevertheless, in this case, you can rely on the security of the access list.

Dynamic NAT has these disadvantages:

- If the mapped pool has fewer addresses than the real group, you could run out of addresses if the amount of traffic is more than expected.

  Use PAT if this event occurs often, because PAT provides over 64,000 translations using ports of a single address.

- You have to use a large number of routable addresses in the mapped pool; if the destination network requires registered addresses, such as the Internet, you might encounter a shortage of usable addresses.

The advantage of dynamic NAT is that some protocols cannot use PAT. PAT does not work with the following:

- IP protocols that do not have a port to overload, such as GRE version 0.
- Some multimedia applications that have a data stream on one port, the control path on another port, and are not open standard.

See the "When to Use Application Protocol Inspection" section on page 25-2 for more information about NAT and PAT support.

PAT

PAT translates multiple real addresses to a single mapped IP address. Specifically, the security appliance translates the real address and source port (real socket) to the mapped address and a unique port above 1024 (mapped socket). Each connection requires a separate translation, because the source port differs for each connection. For example, 10.1.1.1:1025 requires a separate translation from 10.1.1.1:1026.

After the connection expires, the port translation also expires after 30 seconds of inactivity. The timeout is not configurable. Users on the destination network cannot reliably initiate a connection to a host that uses PAT (even if the connection is allowed by an access list). Not only can you not predict the real or mapped port number of the host, but the security appliance does not create a translation at all unless the translated host is the initiator. See the following "Static NAT" or "Static PAT" sections for reliable access to hosts.

PAT lets you use a single mapped address, thus conserving routable addresses. You can even use the security appliance interface IP address as the PAT address. PAT does not work with some multimedia applications that have a data stream that is different from the control path. See the "When to Use Application Protocol Inspection" section on page 25-2 for more information about NAT and PAT support.

Note For the duration of the translation, a remote host can initiate a connection to the translated host if an access list allows it. Because the port address (both real and mapped) is unpredictable, a connection to the host is unlikely. Nevertheless, in this case, you can rely on the security of the access list. However, policy PAT does not support time-based ACLs.

Static NAT

Static NAT creates a fixed translation of real address(es) to mapped address(es).With dynamic NAT and PAT, each host uses a different address or port for each subsequent translation. Because the mapped address is the same for each consecutive connection with static NAT, and a persistent translation rule exists, static NAT allows hosts on the destination network to initiate traffic to a translated host (if an access list exists that allows it).

The main difference between dynamic NAT and a range of addresses for static NAT is that static NAT allows a remote host to initiate a connection to a translated host (if an access list exists that allows it), while dynamic NAT does not. You also need an equal number of mapped addresses as real addresses with static NAT.

Static PAT

Static PAT is the same as static NAT, except that it lets you specify the protocol (TCP or UDP) and port for the real and mapped addresses.

This feature lets you identify the same mapped address across many different static statements, provided the port is different for each statement. You cannot use the same mapped address for multiple static NAT statements.

For applications that require inspection for secondary channels (for example, FTP and VoIP), the security appliance automatically translates the secondary ports.

For example, if you want to provide a single address for remote users to access FTP, HTTP, and SMTP, but these are all actually different servers on the real network, you can specify static PAT statements for each server that uses the same mapped IP address, but different ports (see Figure 22-8).

Figure 22-8 Static PAT

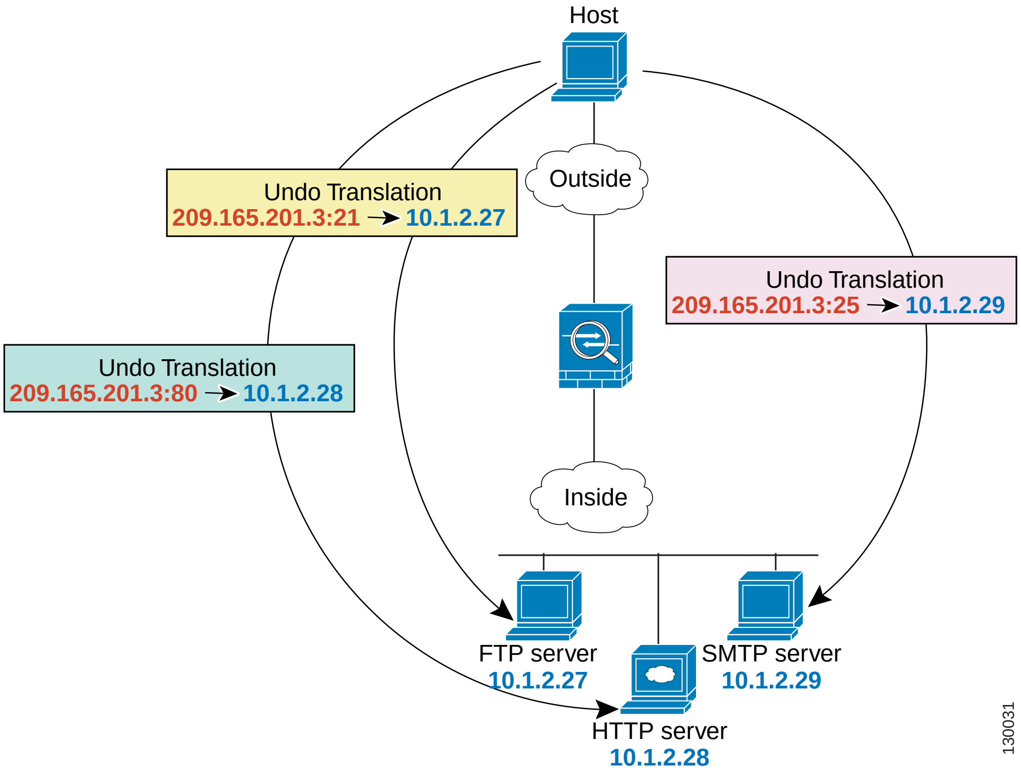

You can also use static PAT to translate a well-known port to a non-standard port or vice versa. For example, if inside web servers use port 8080, you can allow outside users to connect to port 80, and then undo translation to the original port 8080. Similarly, to provide extra security, you can tell web users to connect to non-standard port 6785, and then undo translation to port 80.

Bypassing NAT When NAT Control is Enabled

If you enable NAT control, then inside hosts must match a NAT rule when accessing outside hosts. If you do not want to perform NAT for some hosts, then you can bypass NAT for those hosts or you can disable NAT control. You might want to bypass NAT, for example, if you are using an application that does not support NAT. See the "When to Use Application Protocol Inspection" section on page 25-2 for information about inspection engines that do not support NAT.

You can configure traffic to bypass NAT using one of three methods. All methods achieve compatibility with inspection engines. However, each method offers slightly different capabilities, as follows:

- Identity NAT—When you configure identity NAT (which is similar to dynamic NAT), you do not limit translation for a host on specific interfaces; you must use identity NAT for connections through all interfaces. Therefore, you cannot choose to perform normal translation on real addresses when you access interface A, but use identity NAT when accessing interface B. Regular dynamic NAT, on the other hand, lets you specify a particular interface on which to translate the addresses. Make sure that the real addresses for which you use identity NAT are routable on all networks that are available according to your access lists.

  For identity NAT, even though the mapped address is the same as the real address, you cannot initiate a connection from the outside to the inside (even if the interface access list allows it). Use static identity NAT or NAT exemption for this functionality.
- Static identity NAT—Static identity NAT lets you specify the interface on which you want to allow the real addresses to appear, so you can use identity NAT when you access interface A, and use regular translation when you access interface B. Static identity NAT also lets you use policy NAT, which identifies the real and destination addresses when determining the real addresses to translate (see the "Policy NAT" section on page 22-10 for more information about policy NAT). For example, you can use static identity NAT for an inside address when it accesses the outside interface and the destination is server A, but use a normal translation when accessing the outside server B.
- NAT exemption—NAT exemption allows both translated and remote hosts to initiate connections. Like identity NAT, you do not limit translation for a host on specific interfaces; you must use NAT exemption for connections through all interfaces. However, NAT exemption does let you specify the real and destination addresses when determining the real addresses to translate (similar to policy NAT), so you have greater control using NAT exemption. However unlike policy NAT, NAT exemption does not consider the ports in the access list. NAT exemption also does not let you configure connection limits such as maximum TCP connections.

Policy NAT

Policy NAT lets you identify real addresses for address translation by specifying the source and destination addresses. You can also optionally specify the source and destination ports. Regular NAT can only consider the source addresses, and not the destination. For example, with policy NAT, you can translate the real address to mapped address A when it accesses server A, but translate the real address to mapped address B when it accesses server B.

For applications that require application inspection for secondary channels (for example, FTP and VoIP), the policy specified in the policy NAT rule should include the secondary ports. When the ports cannot be predicted, the policy should specify only the IP addresses for the secondary channel. With this configuration, the security appliance translates the secondary ports.

Figure 22-9 shows a host on the 10.1.2.0/24 network accessing two different servers. When the host accesses the server at 209.165.201.11, the real address is translated to 209.165.202.129. When the host accesses the server at 209.165.200.225, the real address is translated to 209.165.202.130. Consequently, the host appears to be on the same network as the servers, which can help with routing.

Figure 22-9 Policy NAT with Different Destination Addresses

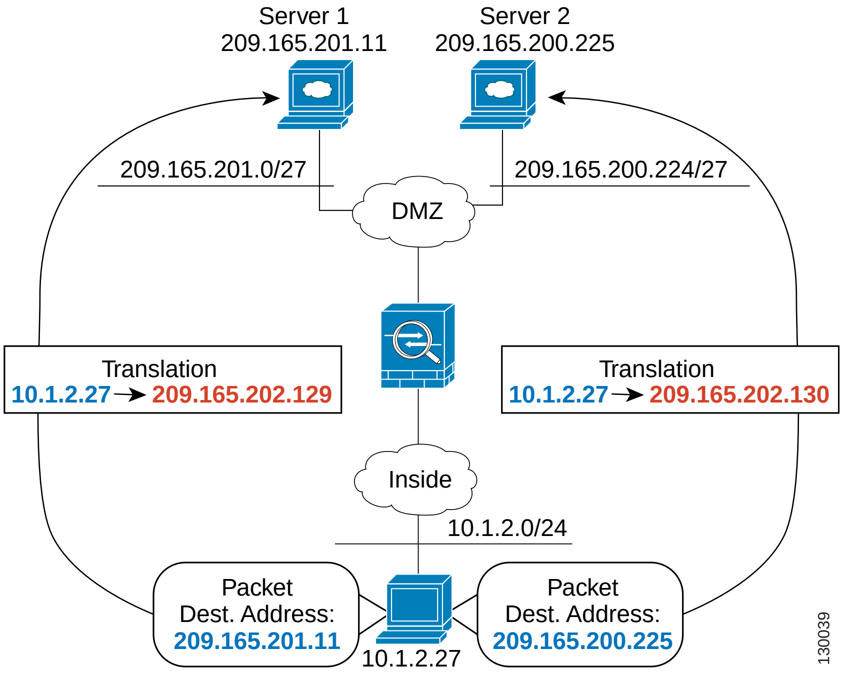

Figure 22-10 shows the use of source and destination ports. The host on the 10.1.2.0/24 network accesses a single host for both web services and Telnet services. When the host accesses the server for web services, the real address is translated to 209.165.202.129. When the host accesses the same server for Telnet services, the real address is translated to 209.165.202.130.

Figure 22-10 Policy NAT with Different Destination Ports

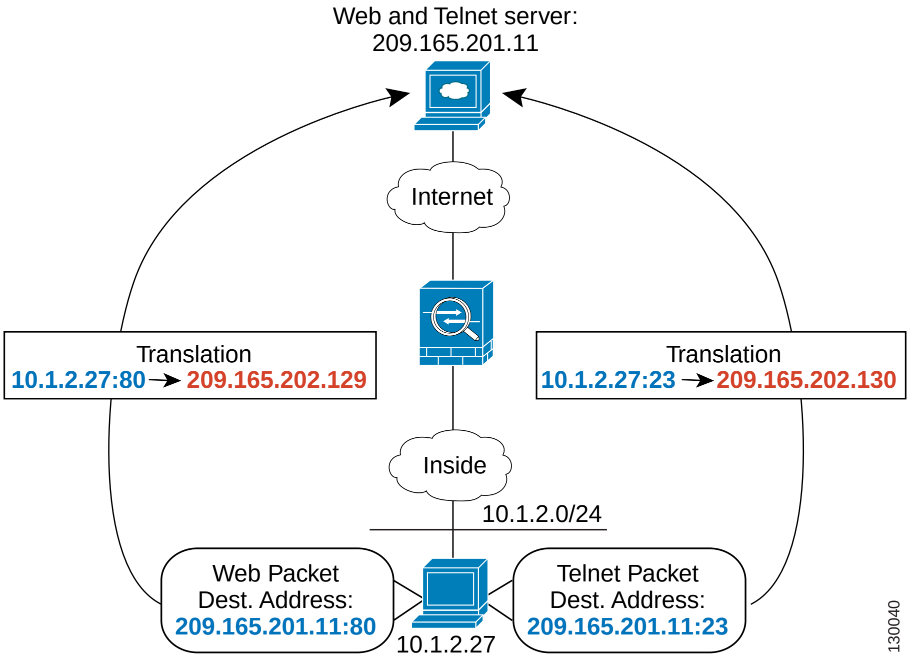

For policy static NAT, both translated and remote hosts can originate traffic. For traffic originated on the translated network, the NAT rule specifies the real addresses and the *destination* addresses, but for traffic originated on the remote network, the rule identifies the real addresses and the *source* addresses of remote hosts who are allowed to connect to the host using this translation.

Figure 22-11 shows a remote host connecting to a translated host. The translated host has a policy static NAT translation that translates the real address only for traffic to and from the 209.165.201.0/27 network. A translation does not exist for the 209.165.200.224/27 network, so the translated host cannot connect to that network, nor can a host on that network connect to the translated host.

Figure 22-11 Policy Static NAT with Destination Address Translation

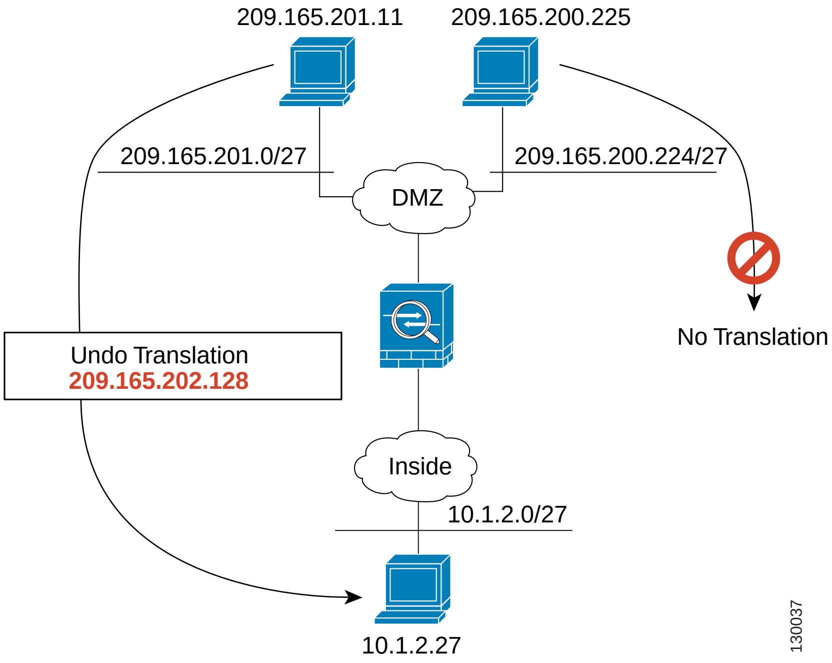

Note Policy NAT does not support SQL*Net, but it is supported by regular NAT. See the "When to Use Application Protocol Inspection" section on page 25-2 for information about NAT support for other protocols.

NAT and Same Security Level Interfaces

NAT is not required between same security level interfaces even if you enable NAT control. You can optionally configure NAT if desired. However, if you configure dynamic NAT when NAT control is enabled, then NAT is required. See the "NAT Control" section on page 22-4 for more information. Also, when you specify a group of IP address(es) for dynamic NAT or PAT on a same security interface, then you must perform NAT on that group of addresses when they access any lower or same security level interface (even when NAT control is not enabled). Traffic identified for static NAT is not affected.

Note The security appliance does not support VoIP inspection engines when you configure NAT on same security interfaces. These inspection engines include Skinny, SIP, and H.323. See the "When to Use Application Protocol Inspection" section on page 25-2 for supported inspection engines.

Order of NAT Rules Used to Match Real Addresses

The security appliance matches real addresses to NAT rules in the following order:

1. NAT exemption—In order, until the first match.
2. Static NAT and Static PAT (regular and policy)—In order, until the first match. Static identity NAT is included in this category.
3. Policy dynamic NAT—In order, until the first match. Overlapping addresses are allowed.
4. Regular dynamic NAT—Best match. Regular identity NAT is included in this category. The order of the NAT rules does not matter; the NAT rule that best matches the real address is used. For example, you can create a general rule to translate all addresses (0.0.0.0) on an interface. If you want to translate a subset of your network (10.1.1.1) to a different address, then you can create a rule to translate only 10.1.1.1. When 10.1.1.1 makes a connection, the specific rule for 10.1.1.1 is used because it matches the real address best. We do not recommend using overlapping rules; they use more memory and can slow the performance of the security appliance.

Mapped Address Guidelines

When you translate the real address to a mapped address, you can use the following mapped addresses:

- Addresses on the same network as the mapped interface.

  If you use addresses on the same network as the mapped interface (through which traffic exits the security appliance), the security appliance uses proxy ARP to answer any requests for mapped addresses, and thus intercepts traffic destined for a real address. This solution simplifies routing, because the security appliance does not have to be the gateway for any additional networks. However, this approach does put a limit on the number of available addresses used for translations.

  For PAT, you can even use the IP address of the mapped interface.

- Addresses on a unique network.

  If you need more addresses than are available on the mapped interface network, you can identify addresses on a different subnet. The security appliance uses proxy ARP to answer any requests for mapped addresses, and thus intercepts traffic destined for a real address. If you use OSPF, and you advertise routes on the mapped interface, then the security appliance advertises the mapped addresses. If the mapped interface is passive (not advertising routes) or you are using static routing, then you need to add a static route on the upstream router that sends traffic destined for the mapped addresses to the security appliance.

DNS and NAT

You might need to configure the security appliance to modify DNS replies by replacing the address in the reply with an address that matches the NAT configuration. You can configure DNS modification when you configure each translation.

For example, a DNS server is accessible from the outside interface. A server, ftp.cisco.com, is on the inside interface. You configure the security appliance to statically translate the ftp.cisco.com real address (10.1.3.14) to a mapped address (209.165.201.10) that is visible on the outside network (see Figure 22-12). In this case, you want to enable DNS reply modification on this static statement so that inside users who have access to ftp.cisco.com using the real address receive the real address from the DNS server, and not the mapped address.

When an inside host sends a DNS request for the address of ftp.cisco.com, the DNS server replies with the mapped address (209.165.201.10). The security appliance refers to the static statement for the inside server and translates the address inside the DNS reply to 10.1.3.14. If you do not enable DNS reply modification, then the inside host attempts to send traffic to 209.165.201.10 instead of accessing ftp.cisco.com directly.

Figure 22-12 DNS Reply Modification

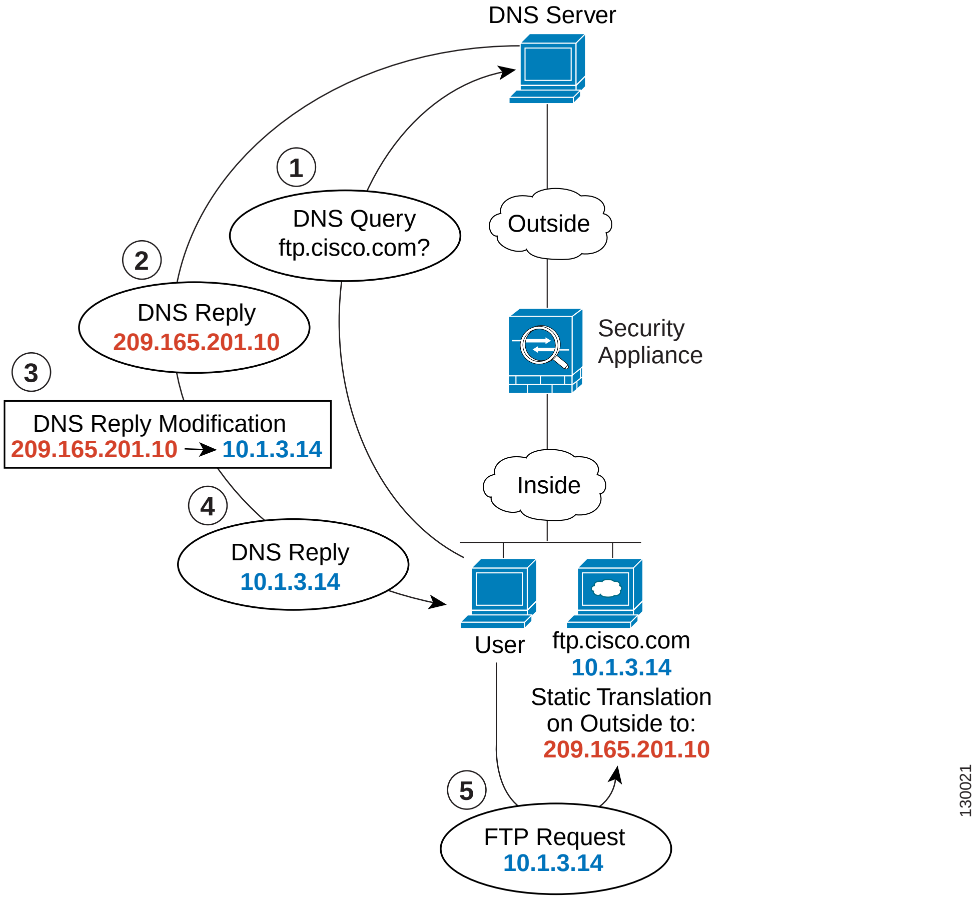

Note If a user on a different network (for example, DMZ) also requests the IP address for ftp.cisco.com from the outside DNS server, then the IP address in the DNS reply is also modified for this user, even though the user is not on the Inside interface referenced by the static rule.

Figure 22-13 shows a web server and DNS server on the outside. The security appliance has a static translation for the outside server. In this case, when an inside user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.20.10. Because you want inside users to use the mapped address for ftp.cisco.com (10.1.2.56) you need to configure DNS reply modification for the static translation.

Figure 22-13 DNS Reply Modification Using Outside NAT

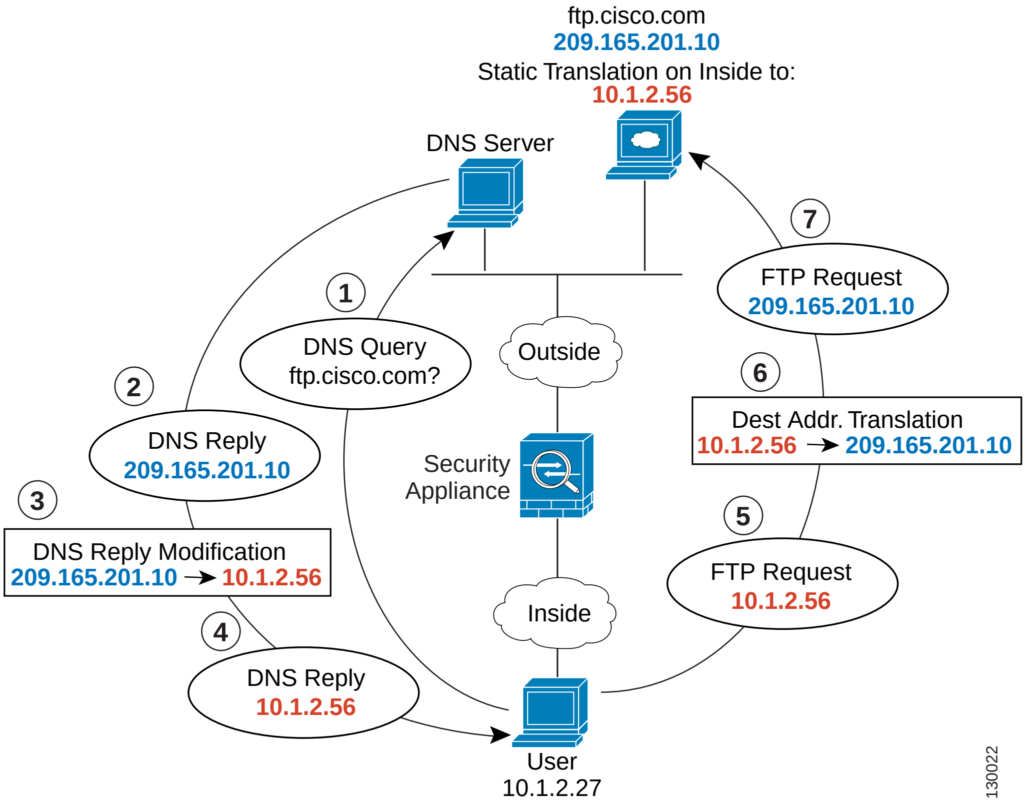

Configuring NAT Control

NAT control requires that packets traversing from an inside interface to an outside interface match a NAT rule. See the "NAT Control" section on page 22-4 for more information.

To enable NAT control, in the Configuration > Firewall > NAT Rules pane, check the **Enable traffic through the firewall without address translation** check box.

Using Dynamic NAT

This section describes how to configure dynamic NAT, including dynamic NAT and PAT, dynamic policy NAT and PAT, and identity NAT.

Policy NAT lets you identify real addresses for address translation by specifying the source and destination addresses. You can also optionally specify the source and destination ports. Regular NAT can only consider the source addresses, and not the destination. See the "Policy NAT" section on page 22-10 for more information.

This section includes the following topics:

- Dynamic NAT Implementation, page 22-16
- Managing Global Pools, page 22-21
- Configuring Dynamic NAT, PAT, or Identity NAT, page 22-22
- Configuring Dynamic Policy NAT or PAT, page 22-24

Dynamic NAT Implementation

This section describes how dynamic NAT is implemented, and includes the following topics:

- Real Addresses and Global Pools Paired Using a Pool ID, page 22-17
- NAT Rules on Different Interfaces with the Same Global Pools, page 22-17
- Global Pools on Different Interfaces with the Same Pool ID, page 22-18
- Multiple NAT Rules with Different Global Pools on the Same Interface, page 22-18
- Multiple Addresses in the Same Global Pool, page 22-19
- Outside NAT, page 22-20
- Real Addresses in a NAT Rule Must be Translated on All Lower or Same Security Interfaces, page 22-21

Real Addresses and Global Pools Paired Using a Pool ID

In a dynamic NAT rule, you specify real addresses and then pair them with a global pool of addresses to which the real addresses are mapped when they exit another interface (in the case of PAT, this is one address, and in the case of identity NAT, this is the same as the real address). Each global pool is assigned a pool ID.

NAT Rules on Different Interfaces with the Same Global Pools

You can create a NAT rule for each interface using the same global address pool. For example, you can configure NAT rules for Inside and DMZ interfaces, both using global pool 1 on the outside interface. Traffic from the Inside interface and the DMZ interface share a mapped pool or a PAT address when exiting the Outside interface (see Figure 22-14).

Figure 22-14 NAT Rules on Multiple Interfaces Using the Same Global Pool

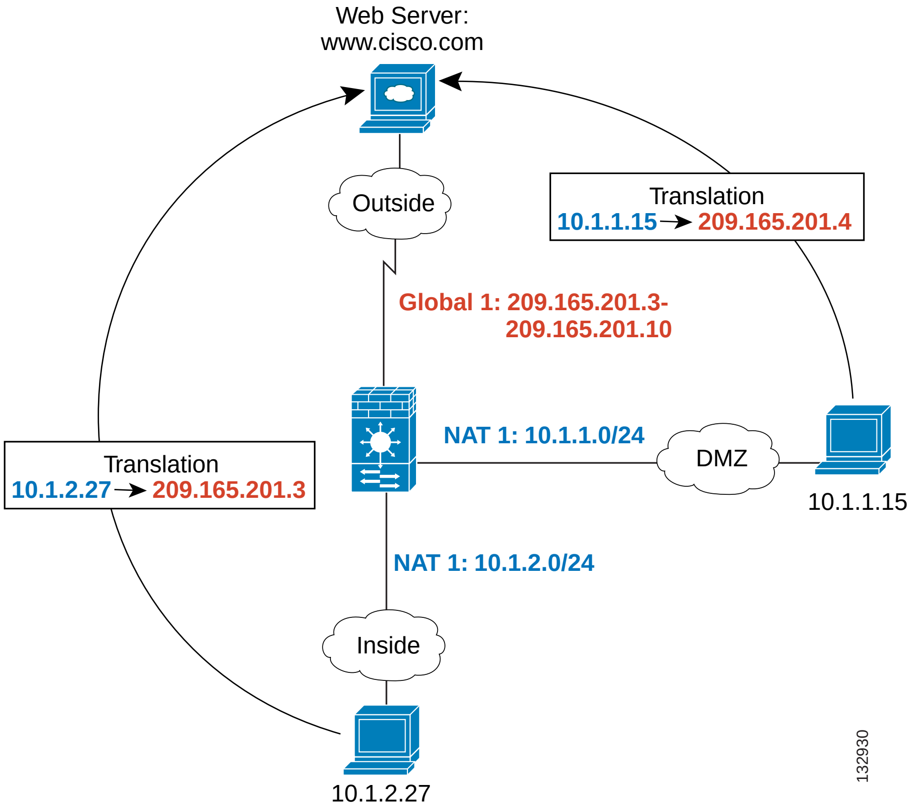

Global Pools on Different Interfaces with the Same Pool ID

You can create a global pool for each interface using the same pool ID. If you create a global pool for the Outside and DMZ interfaces on ID 1, then a single NAT rule associated with ID 1 identifies traffic to be translated when going to both the Outside and the DMZ interfaces. Similarly, if you create a NAT rule for the DMZ interface on ID 1, then all global pools on ID 1 are also used for DMZ traffic. (See Figure 22-15).

Figure 22-15 NAT Rules and Global Pools using the Same ID on Multiple Interfaces

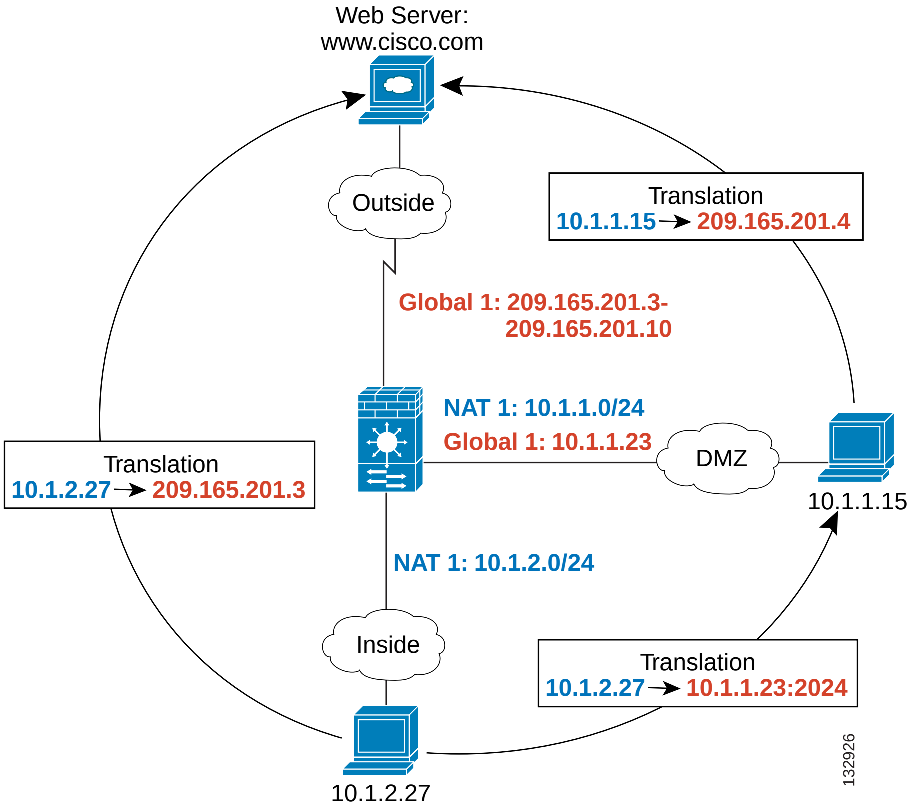

Multiple NAT Rules with Different Global Pools on the Same Interface

You can identify different sets of real addresses to have different mapped addresses. For example, on the Inside interface, you can have two NAT rules on two different pool IDs. On the Outside interface, you configure two global pools for these two IDs. Then, when traffic from Inside network A exits the Outside interface, the IP addresses are translated to pool 1 addresses; while traffic from Inside network B are translated to pool 2 addresses (see Figure 22-16). If you use policy NAT, you can specify the same real addresses for multiple NAT rules, as long as the destination addresses and ports are unique in each access list.

Figure 22-16 Different NAT IDs

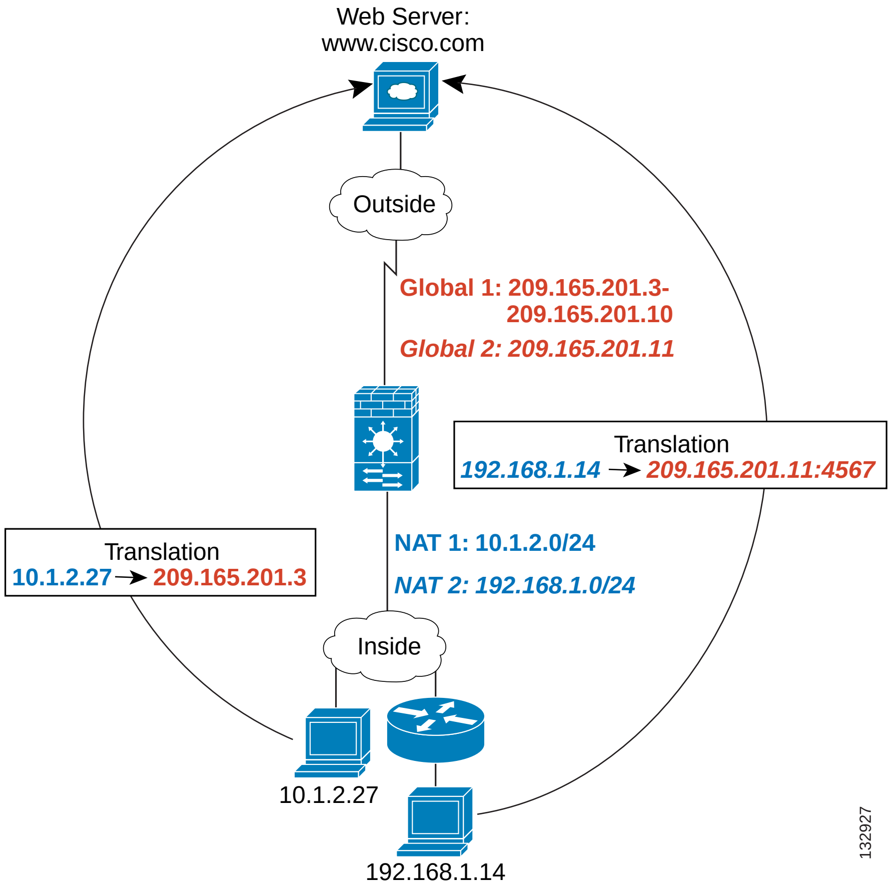

Multiple Addresses in the Same Global Pool

You can have multiple addresses in the same global pool; the security appliance uses the dynamic NAT ranges of addresses first, in the order they are in the configuration, and then uses the PAT single addresses in order. You might want to add both a range of addresses and a PAT address if you need to use dynamic NAT for a particular application, but want to have a backup PAT rule in case all the dynamic NAT addresses are depleted. Similarly, you might want two PAT addresses in the pool if you need more than the approximately 64,000 PAT sessions that a single PAT mapped address supports (see Figure 22-17).

Figure 22-17 NAT and PAT Together

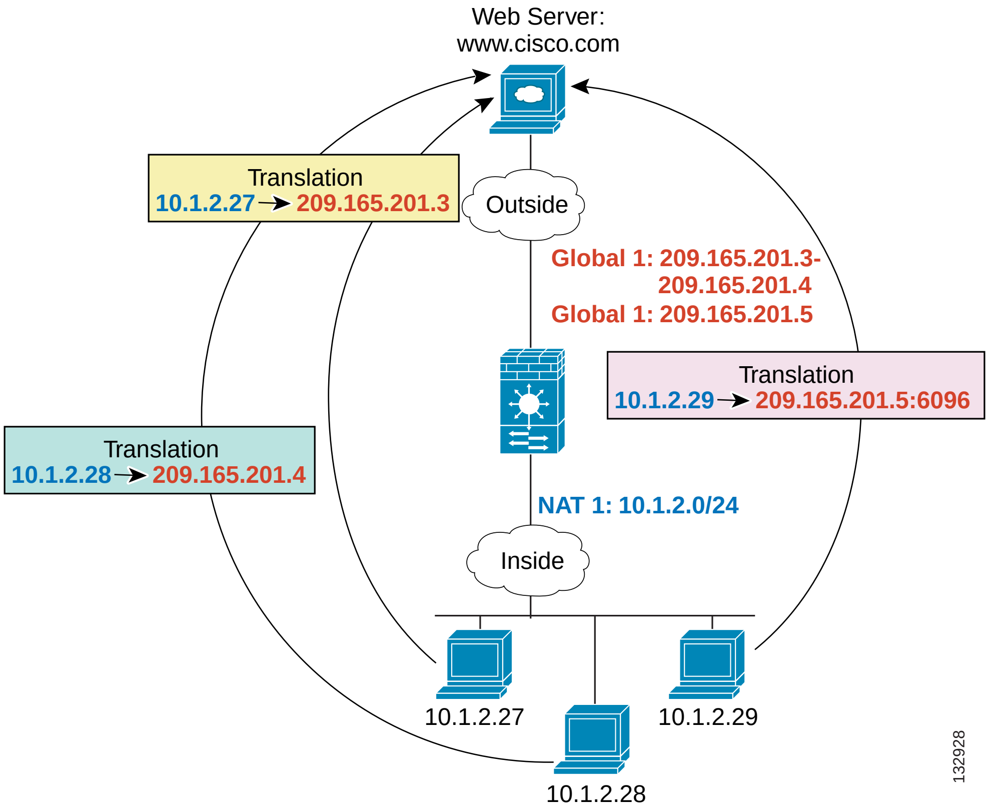

Outside NAT

If a NAT rule translates addresses from an outside interface to an inside interface, then the rule is an outside NAT rule, and you need to specify that it translates inbound traffic. If you also want to translate the same traffic when it accesses a lower security interface (for example, traffic on a DMZ is translated when accessing the Inside and the Outside interfaces), then you can create a second NAT rule using the same NAT ID (see Figure 22-18), but specifying outbound. Note that for outside NAT (DMZ interface to Inside interface), the inside host uses a static rule to allow outside access, so both the source and destination addresses are translated.

Figure 22-18 Outside NAT and Inside NAT Combined

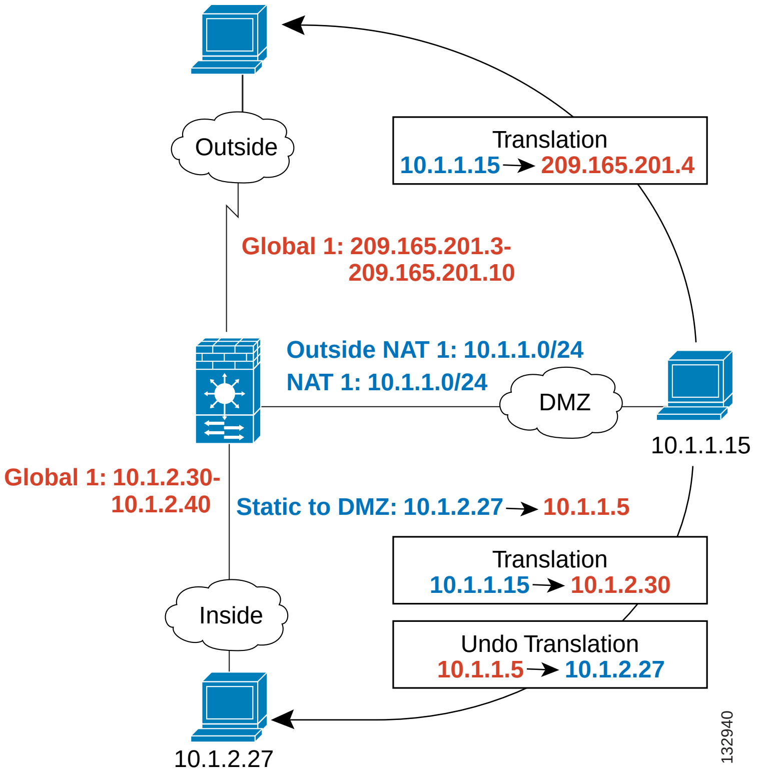

Real Addresses in a NAT Rule Must be Translated on All Lower or Same Security Interfaces

When you create a NAT rule for a group of IP addresses, then you must perform NAT on that group of addresses when they access any lower or same security level interface; you must create a global pool with the same pool ID on each interface, or use a static rule. NAT is not required for that group when it accesses a higher security interface. If you create an outside NAT rule, then the NAT requirements preceding come into effect for that group of addresses when they access all higher security interfaces. Traffic identified by a static rule is not affected.

Managing Global Pools

Dynamic NAT uses global pools for translation. For information about how global pools work, see the "Dynamic NAT Implementation" section on page 22-16.

To manage a global pool, perform the following steps:

Step 1 In the Configuration > Firewall > Objects > Global Pools pane, click **Add** to add a new pool, or select a pool, and click **Edit**.

You can also manage global pools from the Add/Edit Dynamic NAT Rule dialog box by clicking **Manage**.

The Add/Edit Global Address Pool dialog box appears.

Step 2 For a new pool, from the Interface drop-down list, choose the interface where you want to use the mapped IP addresses.

Step 3 For a new pool, in the Pool ID field, enter a number between 1 and 2147483647. Do not enter a pool ID that is already in use, or your configuration will be rejected.

Step 4 In the IP Addresses to Add area, click **Range**, **Port Address Translation (PAT)**, or **PAT Address Translation (PAT) Using IP Address of the interface**.

If you specify a range of addresses, the security appliance performs dynamic NAT. If you specify a subnet mask in the Netmask field, the value specifies the subnet mask assigned to the mapped address when it is assigned to a host. If you do not specify a mask, then the default mask for the address class is used.

Step 5 Click **Add** to add the addresses to the Addresses Pool pane.

Step 6 (Optional) You can add multiple addresses to the global pool. If you want to add a PAT address after you configure a dynamic range, for example, then complete the value for PAT and click **Add** again. See the "Multiple Addresses in the Same Global Pool" section on page 22-19 for information about using multiple addresses on the same pool ID for an interface.

Step 7 Click **OK**.

Configuring Dynamic NAT, PAT, or Identity NAT

To configure a dynamic NAT, PAT, or identity NAT rule, perform the following steps.

Step 1 In the Configuration > Firewall > NAT Rules pane, choose **Add > Add Dynamic NAT Rule**.

The Add Dynamic NAT Rule dialog box appears.

Step 2 In the Original area, from the Interface drop-down list, choose the interface that is connected to the hosts with real addresses that you want to translate.

Step 3 Enter the real addresses in the Source field, or click the **...** button to select an IP address that you already defined in ASDM.

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

Step 4 To choose a global pool, use one of the following options:

- Select an already-defined global pool.

  If the pool includes a range of addresses, then the security appliance performs dynamic NAT. If the pool includes a single address, then the security appliance performs dynamic PAT. If a pool includes both ranges and single addresses, then the ranges are used in order, and then the PAT addresses are used in order. See the "Multiple Addresses in the Same Global Pool" section on page 22-19 for more information.

  Pools are identified by a pool ID. If multiple global pools on different interfaces share the same pool ID, then they are grouped. If you choose a multi-interface pool ID, then traffic is translated as specified when it accesses any of the interfaces in the pool. For more information about pool IDs, see the "Dynamic NAT Implementation" section on page 22-16.

- Create a new global pool or edit an existing pool by clicking **Manage**. See the "Managing Global Pools" section on page 22-21.

- Choose identity NAT by selecting **global pool 0**.

Step 5 (Optional) To enable translation of addresses inside DNS replies, expand the **Connection Settings** area, and check the **Translate the DNS replies that match the translation rule** check box.

If your NAT rule includes the real address of a host that has an entry in a DNS server, and the DNS server is on a different interface from a client, then the client and the DNS server need different addresses for the host; one needs the mapped address and one needs the real address. This option rewrites the address in the DNS reply to the client. The mapped host needs to be on the same interface as either the client or the DNS server. Typically, hosts that need to allow access from other interfaces use a static translation, so this option is more likely to be used with a static rule. See the "DNS and NAT" section on page 22-13 for more information.

Step 6 (Optional) To enable connection settings, expand the **Connection Settings** area, and set one or more of the following options:

Note You can also set these values using a security policy rule (see the "Enabling Connection Limits, TCP Normalization , and TCP State Bypass" section on page 28-10). If you set them in both places, 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.

- **Randomize sequence number**—With this check box checked (the default), the security appliance randomizes the sequence number of TCP packets. Each TCP connection has two ISNs: one generated by the client and one generated by the server. The security appliance randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

  Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

  TCP initial sequence number randomization can be disabled if required. For example:

  - If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.
  - If you use eBGP multi-hop through the security appliance, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.
  - You use a WAAS device that requires the security appliance not to randomize the sequence numbers of connections.
- **Maximum TCP Connections**—Specifies the maximum number of TCP connections, between 0 and 65,535. If this value is set to 0, the number of connections is unlimited.
- **Maximum UDP Connections**—Specifies the maximum number of UDP connections, between 0 and 65,535. If this value is set to 0, the number of connections is unlimited.
- **Maximum 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.

Step 7 Click **OK**.

Configuring Dynamic Policy NAT or PAT

To configure dynamic policy NAT or PAT, perform the following steps:

Step 1 In the Configuration > Firewall > NAT Rules pane, choose **Add > Advanced > Add Dynamic Policy NAT Rule**.

The Add Dynamic Policy NAT Rule dialog box appears.

Step 2 In the Original area, from the Interface drop-down list, choose the interface that is connected to the hosts with real addresses that you want to translate.

Step 3 Enter the real addresses in the Source field, or click the **...** button to choose an IP address that you already defined in ASDM.

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

Separate multiple real addresses by a comma.

Step 4 Enter the destination addresses in the Destination field, or click the **...** button to choose an IP address that you already defined in ASDM.

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

Separate multiple destination addresses by a comma.

By default, the field shows **any**, which allows any destination address.

Step 5 To choose a global pool, use one of the following options:

- Choose an already-defined global pool.

  If the pool includes a range of addresses, then the security appliance performs dynamic NAT. If the pool includes a single address, then the security appliance performs dynamic PAT. If a pool includes both ranges and single addresses, then the ranges are used in order, and then the PAT addresses are used in order. See the "Multiple Addresses in the Same Global Pool" section on page 22-19 for more information.

  Pools are identified by a pool ID. If multiple global pools on different interfaces share the same pool ID, then they are grouped. If you choose a multi-interface pool ID, then traffic is translated as specified when it accesses any of the interfaces in the pool. For more information about pool IDs, see the "Dynamic NAT Implementation" section on page 22-16.

- Create a new global pool or edit an existing pool by clicking **Manage**. See the "Managing Global Pools" section on page 22-21.
- Choose identity NAT by choosing global pool 0.

Step 6 (Optional) Enter a description in the Description field.

Step 7 (Optional) To enable translation of addresses inside DNS replies, expand the **Connection Settings** area, and check the **Translate the DNS replies that match the translation rule** check box.

If your NAT rule includes the real address of a host that has an entry in a DNS server, and the DNS server is on a different interface from a client, then the client and the DNS server need different addresses for the host; one needs the mapped address and one needs the real address. This option rewrites the address in the DNS reply to the client. The mapped host needs to be on the same interface as either the client or the DNS server. Typically, hosts that need to allow access from other interfaces use a static translation, so this option is more likely to be used with a static rule. See the "DNS and NAT" section on page 22-13 for more information.

Step 8 (Optional) To enable connection settings, expand the **Connection Settings** area, and set one or more of the following options:

Note You can also set these values using a security policy rule. To set the number of rate intervals maintained for host statistics, on the Configuration > Firewall > Threat Detection > Scanning Threat Statistics area, choose **1**, **2**, or **3** from the User can specify the number of rate for Threat Detection Host drop-down list. Because host statistics use a lot of memory, reducing the number of rate intervals from the default of 3 reduces the memory usage. By default, the Firewall Dashboard Tab shows information for three rate intervals, for example, for the last 1 hour, 8 hours, and 24 hours. If you set this keyword to **1**, then only the shortest rate interval statistics are maintained. If you set the value to **2**, then the two shortest intervals are maintained. If you set them in both places, 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.

- **Randomize sequence number**—With this check box checked (the default), the security appliance randomizes the sequence number of TCP packets. Each TCP connection has two ISNs: one generated by the client and one generated by the server. The security appliance randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

  Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

  TCP initial sequence number randomization can be disabled if required. For example:

  – If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.

  – If you use eBGP multi-hop through the security appliance, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.

  – You use a WAAS device that requires the security appliance not to randomize the sequence numbers of connections.

- **Maximum TCP Connections**—Specifies the maximum number of TCP connections, between 0 and 65,535. If this value is set to 0, the number of connections is unlimited.
- **Maximum UDP Connections**—Specifies the maximum number of UDP connections, between 0 and 65,535. If this value is set to 0, the number of connections is unlimited.
- **Maximum 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.

Step 9 Click **OK**.

Using Static NAT

This section describes how to configure a static translation, using regular or policy static NAT, PAT, or identity NAT.

For more information about static NAT, see the "Static NAT" section on page 22-8.

Policy NAT lets you identify real addresses for address translation by specifying the source and destination addresses. You can also optionally specify the source and destination ports. Regular NAT can only consider the source addresses, and not the destination. See the "Policy NAT" section on page 22-10 for more information.

Static PAT lets you translate the real IP address to a mapped IP address, as well as the real port to a mapped port. You can choose to translate the real port to the same port, which lets you translate only specific types of traffic, or you can take it further by translating to a different port. For applications that require application inspection for secondary channels (for example, FTP and VoIP), the security appliance automatically translates the secondary ports. For more information about static PAT, see the "Static PAT" section on page 22-9.

You cannot use the same real or mapped address in multiple static rules between the same two interfaces unless you use static PAT. Do not use a mapped address in the static rule that is also defined in a global pool for the same mapped interface.

Static identity NAT translates the real IP address to the same IP address.

This section includes the following topics:

- Configuring Static NAT, PAT, or Identity NAT, page 22-26
- Configuring Static Policy NAT, PAT, or Identity NAT, page 22-28

Configuring Static NAT, PAT, or Identity NAT

To configure static NAT, PAT, or identity NAT, perform the following steps:

Step 1 In the Configuration > Firewall > NAT Rules pane, choose **Add > Add Static NAT Rule**.

The Add Static NAT Rule dialog box appears.

Step 2 In the Original area, from the Interface drop-down list, choose the interface that is connected to the hosts with real addresses that you want to translate.

Step 3 Enter the real addresses in the Source field, or click the **...** button to choose an IP address that you already defined in ASDM.

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

Step 4 In the Translated area, from the Interface drop-down list, choose the interface where you want to use the mapped addresses.

Step 5 Specify the mapped IP address by clicking one of the following:

- **Use IP Address**

  Enter the IP address or click the **...** button to choose an IP address that you already defined in ASDM.

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

- **Use Interface IP Address**

The real and mapped addresses must have the same subnet mask.

Note For identity NAT, enter the same IP address in the Original and Translated fields.

Step 6 (Optional) To use static PAT, check **Enable Port Address Translation (PAT)**.

a. For the Protocol, click **TCP** or **UDP**.

b. In the Original Port field, enter the real port number.

c. In the Translated Port field, enter the mapped port number.

Step 7 (Optional) To enable translation of addresses inside DNS replies, expand the **Connection Settings** area, and check the **Translate the DNS replies that match the translation rule** check box.

If your NAT rule includes the real address of a host that has an entry in a DNS server, and the DNS server is on a different interface from a client, then the client and the DNS server need different addresses for the host; one needs the mapped address and one needs the real address. This option rewrites the address in the DNS reply to the client. The mapped host needs to be on the same interface as either the client or the DNS server. See the "DNS and NAT" section on page 22-13 for more information.

Step 8 (Optional) To enable connection settings, expand the **Connection Settings** area, and set one or more of the following options:

Note You can also set these values using a security policy rule. To set the number of rate intervals maintained for host statistics, on the Configuration > Firewall > Threat Detection > Scanning Threat Statistics area, choose **1**, **2**, or **3** from the User can specify the number of rate for Threat Detection Host drop-down list. Because host statistics use a lot of memory, reducing the number of rate intervals from the default of 3 reduces the memory usage. By default, the Firewall Dashboard Tab shows information for three rate intervals, for example, for the last 1 hour, 8 hours, and 24 hours. If you set this keyword to **1**, then only the shortest rate interval statistics are maintained. If you set the value to **2**, then the two shortest intervals are maintained. If you set them in both places, 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.

- **Randomize sequence number**—With this check box checked (the default), the security appliance randomizes the sequence number of TCP packets. Each TCP connection has two ISNs: one generated by the client and one generated by the server. The security appliance randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

  Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

  TCP initial sequence number randomization can be disabled if required. For example:

  - If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.
  - If you use eBGP multi-hop through the security appliance, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.
  - You use a WAAS device that requires the security appliance not to randomize the sequence numbers of connections.

- **Maximum TCP Connections**—Specifies the maximum number of TCP connections, between 0 and 65,535. If this value is set to 0, the number of connections is unlimited.

- **Maximum UDP Connections**—Specifies the maximum number of UDP connections, between 0 and 65,535. If this value is set to 0, the number of connections is unlimited.
- **Maximum 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.

Step 9 Click **OK**.

Configuring Static Policy NAT, PAT, or Identity NAT

To configure static policy NAT, PAT, or identity NAT, perform the following steps:

Step 1 In the Configuration > Firewall > NAT Rules pane, choose **Add > Advanced > Add Static Policy NAT Rule**.

The Add Static Policy NAT Rule dialog box appears.

Step 2 In the Original area, from the Interface drop-down list, choose the interface that is connected to the hosts with real addresses that you want to translate.

Step 3 Enter the real addresses in the Source field, or click the **...** button to choose an IP address that you already defined in ASDM.

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

Step 4 Enter the destination addresses in the Destination field, or click the **...** button to choose an IP address that you already defined in ASDM.

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

Separate multiple destination addresses by a comma.

By default, the field shows **any**, which allows any destination address.

Step 5 In the Translated area, from the Interface drop-down list, choose the interface where you want to use the mapped addresses.

Step 6 Specify the mapped IP address by clicking one of the following:

- **Use IP Address**

  Enter the IP address or click the **...** button to choose an IP address that you already defined in ASDM.

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

- **Use Interface IP Address**

The real and mapped addresses must have the same subnet mask.

Step 7 (Optional) To use static PAT, check **Enable Port Address Translation (PAT)**.

a. For the Protocol, click **TCP** or **UDP**.

b. In the Original Port field, enter the real port number.

c. In the Translated Port field, enter the mapped port number.

Step 8 (Optional) Enter a description in the Description field.

Step 9 (Optional) To enable translation of addresses inside DNS replies, expand the **Connection Settings** area, and check the **Translate the DNS replies that match the translation rule** check box.

If your NAT rule includes the real address of a host that has an entry in a DNS server, and the DNS server is on a different interface from a client, then the client and the DNS server need different addresses for the host; one needs the mapped address and one needs the real address. This option rewrites the address in the DNS reply to the client. The mapped host needs to be on the same interface as either the client or the DNS server. See the "DNS and NAT" section on page 22-13 for more information.

Step 10 (Optional) To enable connection settings, expand the **Connection Settings** area, and set one or more of the following options:

Note You can also set these values using a security policy rule. To set the number of rate intervals maintained for host statistics, on the Configuration > Firewall > Threat Detection > Scanning Threat Statistics area, choose **1**, **2**, or **3** from the User can specify the number of rate for Threat Detection Host drop-down list. Because host statistics use a lot of memory, reducing the number of rate intervals from the default of 3 reduces the memory usage. By default, the the Firewall Dashboard Tab shows information for three rate intervals, for example, for the last 1 hour, 8 hours, and 24 hours. If you set this keyword to **1**, then only the shortest rate interval statistics are maintained. If you set the value to **2**, then the two shortest intervals are maintained. If you set them in both places, 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.

- **Randomize sequence number**—With this check box checked (the default), the security appliance randomizes the sequence number of TCP packets. Each TCP connection has two ISNs: one generated by the client and one generated by the server. The security appliance randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

  Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

  TCP initial sequence number randomization can be disabled if required. For example:

  - If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.
  - If you use eBGP multi-hop through the security appliance, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.
  - You use a WAAS device that requires the security appliance not to randomize the sequence numbers of connections.
- **Maximum TCP Connections**—Specifies the maximum number of TCP connections, between 0 and 65,535. If this value is set to 0, the number of connections is unlimited.
- **Maximum UDP Connections**—Specifies the maximum number of UDP connections, between 0 and 65,535. If this value is set to 0, the number of connections is unlimited.

- **Maximum 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.

Step 11 Click **OK**.

Using NAT Exemption

NAT exemption exempts addresses from translation and allows both real and remote hosts to originate connections. NAT exemption lets you specify the real and destination addresses when determining the real traffic to exempt (similar to policy NAT), so you have greater control using NAT exemption than dynamic identity NAT. However unlike policy NAT, NAT exemption does not consider the ports. Use static policy identity NAT to consider ports.

For more information about NAT exemption, see the "Bypassing NAT When NAT Control is Enabled" section on page 22-10.

To configure NAT exemption, perform the following steps:

Step 1 In the Configuration > Firewall > NAT Rules pane, choose **Add > Add NAT Exempt Rule**.

The Add NAT Exempt Rule dialog box appears.

Step 2 Click **Action: Exempt**.

Step 3 In the Original area, from the Interface drop-down list, choose the interface that is connected to the hosts with real addresses that you want to exempt.

Step 4 Enter the real addresses in the Source field, or click the **...** button to choose an IP address that you already defined in ASDM.

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

Note You can later specify addresses that you do not want to exempt. For example, you can specify a subnet to exempt such as 10.1.1.0/24, but if you want to translate 10.1.1.50, then you can create a separate rule for that address that removes the exemption.

Separate multiple real addresses by a comma.

Step 5 Enter the destination addresses in the Destination field, or click the **...** button to choose an IP address that you already defined in ASDM.

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

Separate multiple destination addresses by a comma.

By default, the field shows **any**, which allows any destination address.

Step 6 In the NAT Exempt Direction area, choose whether you want to exempt traffic going to lower security interfaces (the default) or to higher security interfaces by clicking the appropriate radio button.

Step 7 (Optional) Enter a description in the Description field.

Step 8 Click **OK**.

Step 9 (Optional) If you do not want to exempt some addresses that were included in your NAT exempt rule, then create another rule to remove the exemption. Right-click the existing NAT Exempt rule, and choose **Insert**.

The Add NAT Exempt Rule dialog box appears.

a. Click **Action: Do not exempt**.

b. Complete Steps 3 through 8 to complete the rule.

The No Exempt rule is added before the Exempt rule. The order of Exempt and No Exempt rules is important. When the security appliance decides whether to exempt a packet, the security appliance tests the packet against each NAT exempt and No Exempt rule in the order in which the rules are listed. After a match is found, no more rules are checked.

CHAPTER 23

Configuring Service Policy Rules

This chapter describes how to enable service policy rules. Service policies provide a consistent and flexible way to configure security appliance features. For example, you can use a service policy to create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications.

This chapter includes the following sections:

Information About Service Policies

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

Service Policy Supported Features

Features can be applied to through traffic or to management traffic. This section includes the following topics:

- "Supported Features for Through Traffic" section on page 23-2

- "Supported Features for Management Traffic" section on page 23-2

Supported Features for Through Traffic

Table 23-1 lists the features supported by Modular Policy Framework.

Table 23-1 Modular Policy Framework Features

Feature	See:
Application inspection (multiple types)	Chapter 25, "Configuring Application Layer Protocol Inspection."
CSC	Chapter 31, "Configuring Trend Micro Content Security."
IPS	Chapter 30, "Configuring the IPS Application on the AIP SSM and SSC."
NetFlow Secure Event Logging filtering	"Using NetFlow" section on page 19-18
QoS input and output policing	Chapter 26, "Configuring QoS."
QoS standard priority queue	Chapter 26, "Configuring QoS."
QoS traffic shaping, hierarchical priority queue	Chapter 26, "Configuring QoS."
TCP and UDP connection limits	"Configuring Connection Settings" section on page 28-6
TCP and UDP connection timeouts	"Configuring Connection Settings" section on page 28-6
TCP sequence number randomization	"Configuring Connection Settings" section on page 28-6
TCP normalization	"Configuring Connection Settings" section on page 28-6
TCP state bypass	"Configuring Connection Settings" section on page 28-6

Supported Features for Management Traffic

Table 23-2 lists the features supported by Modular Policy Framework for management traffic.

Table 23-2 Modular Policy Framework Features for Management Traffic

Feature	See:
Application inspection for RADIUS accounting traffic	"RADIUS Accounting Inspection" section on page 25-22
TCP and UDP connection limits	"Configuring Connection Settings" section on page 28-6

Information About Configuring a Service Policy

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

1. Identify the interface to which you want to apply the rule, or identify the global policy.
2. Identify the traffic to which you want to apply actions. You can identify Layer 3 and 4 through traffic.
3. Apply actions to the traffic class. You can apply multiple actions for each traffic class.

Feature Directionality

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

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

For features that are applied unidirectionally, for example QoS priority queue, only traffic that enters (or exits, depending on the feature) the interface to which you apply the policy map is affected. See Table 23-3 for the directionality of each feature.

Table 23-3 Feature Directionality

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

Feature Matching Within a Policy Map

See the following information for how a packet matches class maps in a policy map:

1. A packet can match only one class map in the policy map for each feature type.
2. When the packet matches a class map for a feature type, the security appliance does not attempt to match it to any subsequent class maps for that feature type.
3. If the packet matches a subsequent class map for a different feature type, however, then the security appliance also applies the actions for the subsequent class map, if supported. See the "Adding a Service Policy Rule for Through Traffic" section on page 23-7 for more information about unsupported combinations.

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

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

Note Application inspection includes multiple inspection types, and each inspection type is a separate feature when you consider the matching guidelines above.

Order in Which Multiple Feature Actions are Applied

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

Note NetFlow Secure Event Logging filtering is order-independent.

Actions are performed in the following order:

1. QoS input policing
2. TCP normalization, TCP and UDP connection limits and timeouts, TCP sequence number randomization, and TCP state bypass.

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

3. CSC
4. Application inspection (multiple types)

   The order of application inspections applied when a class of traffic is classified for multiple inspections is as follows. Only one inspection type can be applied to the same traffic. WAAS inspection is an exception, because it can be applied along with other inspections for the same traffic. See the "Adding a Service Policy Rule for Through Traffic" section on page 23-7 for more information.

   a. CTIQBE
   b. DNS
   c. FTP
   d. GTP
   e. H323
   f. HTTP
   g. ICMP
   h. ICMP error
   i. ILS
   j. MGCP
   k. NetBIOS

- l. PPTP
- m. Sun RPC
- n. RSH
- o. RTSP
- p. SIP
- q. Skinny
- r. SMTP
- s. SNMP
- t. SQL*Net
- u. TFTP
- v. XDMCP
- w. DCERPC
- x. Instant Messaging

Note RADIUS accounting is not listed because it is the only inspection allowed on management traffic. WAAS is not listed because it can be configured along with other inspections for the same traffic.

5. IPS
6. QoS output policing
7. QoS standard priority queue
8. QoS traffic shaping, hierarchical priority queue

Incompatibility of Certain Feature Actions

Some features are not compatible with each other for the same traffic. For example, you cannot configure QoS priority queueing and QoS policing for the same set of traffic. Also, most inspections should not be combined with another inspection, so the security appliance only applies one inspection if you configure multiple inspections for the same traffic. In this case, the feature that is applied is the higher priority feature in the list in the "Order in Which Multiple Feature Actions are Applied" section on page 23-4.

For information about compatibility of each feature, see the chapter or section for your feature.

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

Feature Matching for Multiple Policy Maps

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

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

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

Licensing Requirements for Modular Policy Framework

Model	License Requirement
All models	Base License.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

IPv6 Guidelines

Supports IPv6 for the following features:

- Application inspection for FTP, HTTP, ICMP, SIP, SMTP and IPSec-pass-thru
- IPS
- NetFlow Secure Event Logging filtering
- TCP and UDP connection limits
- TCP and UDP connection timeouts
- TCP sequence number randomization
- TCP normalization
- TCP state bypass

\Service Policy Guidelines

- Interface service policies take precedence over the global service policy for a given feature. For example, if you have a global policy with FTP inspection, and an interface policy with TCP normalization, then both FTP inspection and TCP normalization are applied to the interface. However, if you have a global policy with FTP inspection, and an interface policy with FTP inspection, then only the interface policy FTP inspection is applied to that interface.
- You can only apply one global policy. For example, you cannot create a global policy that includes feature set 1, and a separate global policy that includes feature set 2. All features must be included in a single policy.

Default Settings

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

The default policy includes the following application inspections:

- DNS inspection for the maximum message length of 512 bytes
- FTP
- H323 (H225)
- H323 (RAS)
- RSH
- RTSP
- ESMTP
- SQLnet
- Skinny (SCCP)
- SunRPC
- XDMCP
- SIP
- NetBios
- TFTP

Adding a Service Policy Rule for Through Traffic

See the "Supported Features for Through Traffic" section on page 23-2 for more information. To add a service policy rule for through traffic, perform the following steps:

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

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

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

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

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

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

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

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

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

- **Global - applies to all interfaces**. This option applies the service policy globally to all interfaces. By default, a global policy exists that includes a service policy rule for default application inspection. See the "Default Settings" section on page 23-7 for more information. You can add a rule to the global policy using the wizard.

Step 3 Click **Next**.

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

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

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

  Identify the traffic using one of several criteria:

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

    This option, which is used in the default global policy, is a special shortcut that when used in a rule, ensures that the correct inspection is applied to each packet, based on the destination port of the traffic. For example, when UDP traffic for port 69 reaches the security appliance, then the security appliance applies the TFTP inspection; when TCP traffic for port 21 arrives, then the security appliance applies the FTP inspection. So in this case only, you can configure multiple inspections for the same rule (See the "Adding a Service Policy Rule for Through Traffic" section on page 23-7 for more information about combining actions). Normally, the security appliance does not use the port number to determine the inspection applied, thus giving you the flexibility to apply inspections to non-standard ports, for example.

    See the "Default Inspection Policy" section on page 25-3 for a list of default ports. The security appliance includes a default global policy that matches the default inspection traffic, and applies common inspections to the traffic on all interfaces. Not all applications whose ports are included in the Default Inspection Traffic class are enabled by default in the policy map.

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

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

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

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

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

- **RTP Range**—The class map matches RTP traffic.
- **IP DiffServ CodePoints (DSCP)**—The class matches up to eight DSCP values in the IP header.
- **IP Precedence**—The class map matches up to four precedence values, represented by the TOS byte in the IP header.
- **Any Traffic**—Matches all traffic.

- **Add rule to existing traffic class**. If you already have a service policy rule on the same interface, or you are adding to the global service policy, this option lets you add an ACE to an existing access list. You can add an ACE to any access list that you previously created when you chose the Source and Destination IP Address (uses ACL) option for a service policy rule on this interface. For this traffic class, you can have only one set of rule actions even if you add multiple ACEs. You can add multiple ACEs to the same traffic class by repeating this entire procedure. See the "Managing the Order of Service Policy Rules" section on page 23-15 for information about changing the order of ACEs.
- **Use an existing traffic class.** If you created a traffic class used by a rule on a different interface, you can reuse the traffic class definition for this rule. Note that if you alter the traffic class for one rule, the change is inherited by all rules that use that traffic class. If your configuration includes any **class-map** commands that you entered at the CLI, those traffic class names are also available (although to view the definition of the traffic class, you need to create the rule).
- **Use class default as the traffic class**. This option uses the class-default class, which matches all traffic. The class-default class is created automatically by the security appliance and placed at the end of the policy. If you do not apply any actions to it, it is still created by the security appliance, but for internal purposes only. You can apply actions to this class, if desired, which might be more convenient than creating a new traffic class that matches all traffic. You can only create one rule for this service policy using the class-default class, because each traffic class can only be associated with a single rule per service policy.

Step 5 Click **Next**.

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

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

- Default Inspections—This dialog box is informational only, and shows the applications and the ports that are included in the traffic class.
- Source and Destination Address—This dialog box lets you set the source and destination addresses:

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

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

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

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

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

Separate multiple addresses by a comma.

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

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

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

Separate multiple addresses by a comma.

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

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

By default, the service is IP.

Separate multiple services by a comma.

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

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

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

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

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

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

To add a new time range, click the **...** button. See the "Configuring Time Ranges" section on page 20-15 for more information.

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

- Tunnel Group—Choose a tunnel group from the Tunnel Group drop-down list, or click **New** to add a new tunnel group. See the "IPsec Remote Access Connection Profiles" section on page 36-63 for more information.

  To police each flow, check **Match flow destination IP address**. All traffic going to a unique IP destination address is considered a flow.
- Destination Port—Click **TCP** or **UDP**.

  In the Service field, enter a port number or name, or click **...** to choose one already defined in ASDM.
- RTP Range—Enter an RTP port range, between 2000 and 65534. The maximum number of port sin the range is 16383.
- IP DiffServ CodePoints (DSCP)—In the DSCP Value to Add area, choose a value from the **Select Named DSCP Values** or enter a value in the **Enter DSCP Value (0-63) field**, and click **Add**.

  Add additional values as desired, or remove them using the **Remove** button.
- IP Precedence—From the Available IP Precedence area, choose a value and click **Add**.

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

Step 7 Click **Next**.

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

Step 8 Configure one or more rule actions. See the "Service Policy Supported Features" section on page 23-1 for a list of features.

Step 9 Click **Finish**.

Adding a Service Policy Rule for Management Traffic

You can create a service policy for traffic directed to the security appliance for management purposes. See the "Supported Features for Management Traffic" section on page 23-2 for more information. This section includes the following topics:

- RADIUS Accounting Inspection Overview, page 23-11
- Configuring a Service Policy Rule for Management Traffic, page 23-12

RADIUS Accounting Inspection Overview

One of the well known problems is the over-billing attack in GPRS networks. The over-billing attack can cause consumers anger and frustration by being billed for services that they have not used. In this case, a malicious attacker sets up a connection to a server and obtains an IP address from the SGSN. When the attacker ends the call, the malicious server will still send packets to it, which gets dropped by the GGSN, but the connection from the server remains active. The IP address assigned to the malicious attacker gets released and reassigned to a legitimate user who will then get billed for services that the attacker will use.

RADIUS accounting inspection prevents this type of attack using by ensuring the traffic seen by the GGSN is legitimate. With the RADIUS accounting feature properly configured, the security appliance tears down a connection based on matching the Framed IP attribute in the Radius Accounting Request

Start message with the Radius Accounting Request Stop message. When the Stop message is seen with the matching IP address in the Framed IP attribute, the security appliance looks for all connections with the source matching the IP address.

You have the option to configure a secret pre-shared key with the RADIUS server so the security appliance can validate the message. If the shared secret is not configured, the security appliance does not need to validate the source of the message and will only check that the source IP address is one of the configured addresses allowed to send the RADIUS messages.

Configuring a Service Policy Rule for Management Traffic

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

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

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

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

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

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

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

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

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

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

- **Global - applies to all interfaces**. This option applies the service policy globally to all interfaces. By default, a global policy exists that includes a service policy rule for default application inspection. See the "Default Settings" section on page 23-7 for more information. You can add a rule to the global policy using the wizard.

Step 4 Click **Next**.

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

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

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

  Identify the traffic using one of several criteria:

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

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

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

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

- **Add rule to existing traffic class**. If you already have a service policy rule on the same interface, or you are adding to the global service policy, this option lets you add an ACE to an existing access list. You can add an ACE to any access list that you previously created when you chose the Source and Destination IP Address (uses ACL) option for a service policy rule on this interface. For this traffic class, you can have only one set of rule actions even if you add multiple ACEs. You can add multiple ACEs to the same traffic class by repeating this entire procedure. See the "Managing the Order of Service Policy Rules" section on page 23-15 for information about changing the order of ACEs.
- **Use an existing traffic class.** If you created a traffic class used by a rule on a different interface, you can reuse the traffic class definition for this rule. Note that if you alter the traffic class for one rule, the change is inherited by all rules that use that traffic class. If your configuration includes any **class-map** commands that you entered at the CLI, those traffic class names are also available (although to view the definition of the traffic class, you need to create the rule).

Step 6 Click **Next**.

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

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

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

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

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

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

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

  Separate multiple addresses by a comma.

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

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

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

Separate multiple addresses by a comma.

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

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

By default, the service is IP.

Separate multiple services by a comma.

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

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

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

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

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

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

To add a new time range, click the **...** button. See the "Configuring Time Ranges" section on page 20-15 for more information.

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

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

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

Step 8 Click **Next**.

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

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

See the "RADIUS Accounting Field Descriptions" section on page 23-16 for more information.

Step 10 To configure maximum connections, enter one or more of the following values in the Maximum Connections area:

- **TCP & UDP Connections**—Specifies the maximum number of simultaneous TCP and UDP connections for all clients in the traffic class, up to 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.

Step 11 Click **Finish**.

Managing the Order of Service Policy Rules

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

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

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

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

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

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

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

Step 2 Click the Move Up or Move Down cursor (see Figure 23-1).

Figure 23-1 Moving an ACE

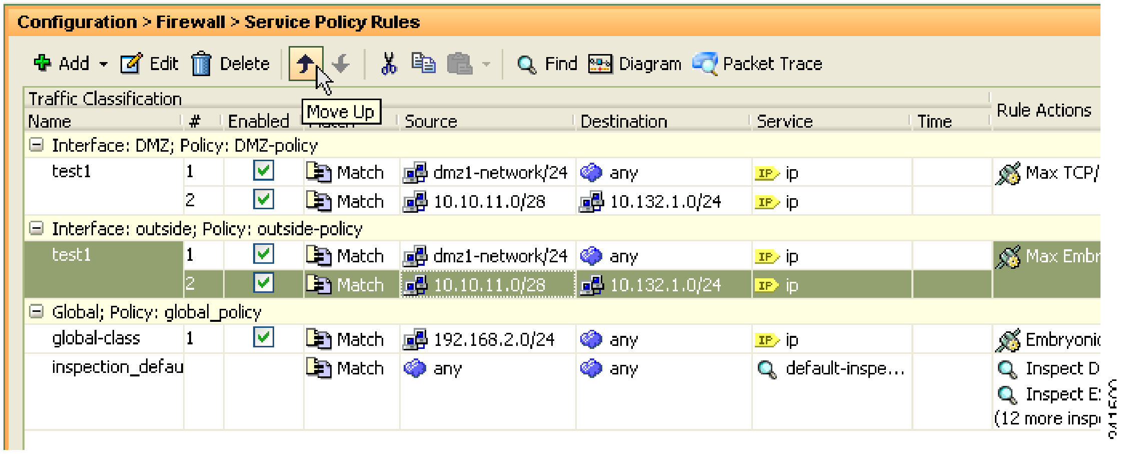

Name	#	Enabled		Source	Destination	Service	Time	Rule Actions
⊟ Interface: DMZ; Policy: DMZ-policy								
test1	1	☑	Match	dmz1-network/24	any	ip		Max TCP/
	2	☑	Match	10.10.11.0/28	10.132.1.0/24	ip		
⊟ Interface: outside; Policy: outside-policy								
test1	1	☑	Match	dmz1-network/24	any	ip		Max Embr
	2	☑	Match	10.10.11.0/28	10.132.1.0/24	ip		
⊟ Global; Policy: global_policy								
global-class	1	☑	Match	192.168.2.0/24	any	ip		Embryonic
inspection_defau			Match	any	any	default-inspe...		Inspect D Inspect E: (12 more insp

> **Note** If you rearrange ACEs in an access list that is used in multiple service policies, then the change is inherited in all service policies.

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

RADIUS Accounting Field Descriptions

This section lists RADIUS accounting field descriptions, and includes the following topics:

- Select RADIUS Accounting Map, page 23-16
- Add RADIUS Accounting Policy Map, page 23-16
- RADIUS Inspect Map, page 23-17
- RADIUS Inspect Map Host, page 23-18
- RADIUS Inspect Map Other, page 23-18

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

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
•	•	•	•	—

RADIUS Inspect Map

The RADIUS pane lets you view previously configured RADIUS application inspection maps. A RADIUS map lets you change the default configuration values used for RADIUS application inspection. You can use a RADIUS map to protect against an overbilling attack.

Fields

- Name—Enter the name of the inspect map, up to 40 characters in length.
- Description—Enter the description of the inspect map, up to 200 characters in length.
- RADIUS Inspect Maps—Table that lists the defined RADIUS inspect maps. The defined inspect maps are also listed in the RADIUS area of the Inspect Maps tree.
- Add—Adds the new RADIUS inspect map to the defined list in the RADIUS Inspect Maps table and to the RADIUS area of the Inspect Maps tree. To configure the new RADIUS map, select the RADIUS entry in Inspect Maps tree.
- Delete—Deletes the application inspection map selected in the RADIUS Inspect Maps table and from the RADIUS area of the Inspect Maps tree.

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

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

The RADIUS Inspect Map Other Parameters pane lets you configure additional parameter settings for the inspect map.

Fields

- Name—Shows the name of the previously configured RADIUS accounting map.
- Description—Enter the description of the RADIUS accounting map, up to 200 characters in length.
- Other Parameters—Lets you configure additional parameters.
  - Send response to the originator of the RADIUS message—Sends a message back to the host from which the RADIUS message was sent.
  - Enforce timeout—Enables the timeout for users.

    Users Timeout—Timeout for the users in the database (hh:mm:ss).
  - Enable detection of GPRS accounting—Enables detection of GPRS accounting. This option is only available when GTP/GPRS license is enabled.
  - Validate Attribute—Attribute information.

    Attribute Number—Specify the attribute number to validate when an Accounting Start is received.

    Add—Adds the entry to the Attribute table.

    Delete—Deletes the entry from the Attribute table.

Modes

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

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

Feature History for Modular Policy Framework

Table 23-4 lists the release history for this feature.

Table 23-4 Feature History for Service Policies

Feature Name	Releases	Feature Information
Modular Policy Framework	7.0(1)	Modular Policy Framework was introduced.
Management class map for use with RADIUS accounting traffic	7.2(1)	The management class map was introduced for use with RADIUS accounting traffic. The following commands were introduced: **class-map type management**, and **inspect radius-accounting**.
Inspection policy maps	7.2(1)	The inspection policy map was introduced. The following command was introduced: **class-map type inspect**.
Regular expressions and policy maps	7.2(1)	Regular expressions and policy maps were introduced to be used under inspection policy maps. The following commands were introduced: **class-map type regex**, **regex**, **match regex**.
Match any for inspection policy maps	8.0(2)	The **match any** keyword was introduced for use with inspection policy maps: traffic can match one or more criteria to match the class map. Formerly, only **match all** was available.
Maximum connections and embryonic connections for management traffic	8.0(2)	The **set connection** command is now available for a Layer 3/4 management class map, for to-the-security appliance management traffic. Only the **conn-max** and **embryonic-conn-max** keywords are available.

CHAPTER 24

Applying AAA for Network Access

This chapter describes how to enable AAA (pronounced “triple A”) for network access.

For information about AAA for management access, see the “Configuring AAA for System Administrators” section on page 18-16.

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 Authentication for Network Access

This section includes the following topics:

Information About Authentication

The security appliance lets you configure network access authentication using AAA servers. This section includes the following topics:

- One-Time Authentication, page 24-2
- Applications Required to Receive an Authentication Challenge, page 24-2
- Security Appliance Authentication Prompts, page 24-2
- Static PAT and HTTP, page 24-3
- Configuring Network Access Authentication, page 24-4

One-Time Authentication

A user at a given IP address only needs to authenticate one time for all rules and types, until the authentication session expires. (See the Configuration > Firewall > Advanced > Global Timeouts pane 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.

Applications Required to Receive an Authentication Challenge

Although you can configure the security appliance to require authentication for network access to any protocol or service, users can authenticate directly with HTTP, HTTPS, Telnet, or FTP only. A user must first authenticate with one of these services before the security appliance allows other traffic requiring authentication.

The authentication ports that the security appliance supports for AAA are fixed:

- Port 21 for FTP
- Port 23 for Telnet
- Port 80 for HTTP
- Port 443 for HTTPS

Security Appliance Authentication Prompts

For Telnet and FTP, the security appliance generates an authentication prompt.

For HTTP, the security appliance uses basic HTTP authentication by default, and provides an authentication prompt. You can optionally configure the security appliance to redirect users to an internal web page where they can enter their username and password (configured on the Configuration > Firewall > AAA Rules > Advanced > AAA Rules Advanced Options dialog box; see the "Enabling the Redirection Method of Authentication for HTTP and HTTPS" section on page 24-5).

For HTTPS, the security appliance generates a custom login screen. You can optionally configure the security appliance to redirect users to an internal web page where they can enter their username and password (configured on the Configuration > Firewall > AAA Rules > Advanced > AAA Rules Advanced Options dialog box; see the "Enabling the Redirection Method of Authentication for HTTP and HTTPS" section on page 24-5).

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; if 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.

After you authenticate correctly, the security appliance redirects you to your original destination. If the destination server also has its own authentication, the user enters another username and password. If you use basic HTTP authentication and need to enter another username and password for the destination server, then you need to configure virtual HTTP (see the Configuration >Firewall > Advanced Options > Virtual Access pane).

Note If you use HTTP authentication, by default the username and password are sent from the client to the security appliance in clear text; in addition, the username and password are sent on to the destination web server as well. See the "Enabling Secure Authentication of Web Clients" section on page 24-5 for information to secure your credentials.

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@patm
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 (@).

Static PAT and HTTP

For HTTP authentication, the security appliance checks real ports when static PAT is configured. If it detects traffic destined for real port 80, regardless of the mapped 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:

```
static (inside,outside) tcp 10.48.66.155 889 192.168.123.10 www netmask 255.255.255.255
```

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, as in the following example:

```
static (inside,outside) tcp 10.48.66.155 889 192.168.123.10 111 netmask 255.255.255.255
```

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.

Configuring Network Access Authentication

To enable network access authentication, perform the following steps. For more information about authentication, see the "Information About Authentication" section on page 24-2.

Step 1 From the Configuration > Firewall > AAA Rules pane, choose **Add > Add Authentication Rule**.

The Add Authentication Rule dialog box appears.

Step 2 From the Interface drop-down list, choose the interface for applying the rule.

Step 3 In the Action field, click one of the following, depending on the implementation:

- **Authenticate**
- **Do not Authenticate**.

Step 4 From the AAA Server Group drop-down list, choose a server group. To add a AAA server to the server group, click **Add Server**. See the "Configuring AAA Server Groups" section on page 16-9 for more information.

If you chose LOCAL for the AAA server group, you can optionally add a new user by clicking **Add User**. See the "Adding a User Account" section on page 16-18 for more information.

Step 5 In the Source field, add the source IP address, or click the ellipsis (...) to choose an IP address already defined in ASDM.

Step 6 In the Destination field, enter the destination IP address, or click the ellipsis (...) to choose an IP address already defined in ASDM.

Step 7 In the Service field, enter an IP service name or number for the destination service, or click ellipsis (...) button to choose a service.

Step 8 (Optional) In the Description field, add a description.

Step 9 (Optional) Click **More Options** to do any of the following:

- To specify a source service for TCP or UDP, enter a TCP or UDP service in the Source Service field.

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

- To make the rule inactive, uncheck **Enable Rule**.

  You may not want to remove a rule, but instead turn it off.

- To set a time range for the rule, from the Time Range drop-down list, choose an existing time range. To add a new time range, click the ellipsis (...). For more information, see Configuring Time Ranges, page 20-15.

Step 10 Click **OK**.

The dialog box closes and the rule appears in the AAA Rules table.

Step 11 Click **Apply**.

The changes are saved to the running configuration.

Enabling the Redirection Method of Authentication for HTTP and HTTPS

This method of authentication enables HTTP(S) listening ports to authenticate network users. When you enable a listening port, the security appliance serves an authentication page for direct connections and, by enabling redirection, for through traffic. This method also prevents the authentication credentials from continuing to the destination server. See the "Security Appliance Authentication Prompts" section on page 24-2 for more information about the redirection method versus the basic method.

To enable a AAA listener, perform the following steps:

Step 1 From the Configuration > Firewall > AAA Rules pane, click **Advanced**.

The AAA Rules Advanced Options dialog box appears.

Step 2 Under Interactive Authentication, click **Add**.

The Add Interactive Authentication Entry dialog box appears.

Step 3 For the Protocol, choose either **HTTP** or **HTTPS**. You can enable both by repeating this procedure and creating two separate rules.

Step 4 From the Interface drop-down list, choose the interface on which you want to enable the listener.

Step 5 From the Port drop-down list, choose the port or enter a number.

This is the port that the security appliance listens on for direct or redirected traffic; the defaults are 80 (HTTP) and 443 (HTTPS). You can use any port number and retain the same functionality, but be sure your direct authentication users know the port number; redirected traffic is sent to the correct port number automatically, but direct authenticators must specify the port number manually.

Step 6 (Optional) Check **Redirect network users for authentication request**.

This option redirects through traffic to an authentication web page served by the security appliance. Without this option, only traffic directed to the security appliance interface can access the authentication web pages.

Note If you enable the redirect option, you cannot also configure static PAT for the same interface where you translate the interface IP address and the same port that is used for the listener; NAT succeeds, but authentication fails.

Step 7 Click **OK**, and then click **OK** to exit the AAA Rules Advanced Options dialog box.

Step 8 Click **Apply**.

Enabling Secure Authentication of Web Clients

If you use HTTP authentication, by default the username and password are sent from the client to the security appliance in clear text; in addition, the username and password are sent on to the destination web server as well. The security appliance provides several methods of securing HTTP authentication, including the following methods:

- Enable the redirection method of authentication for HTTP—See the "Enabling the Redirection Method of Authentication for HTTP and HTTPS" section on page 24-5. This method prevents the authentication credentials from continuing to the destination server.

- Enabling Virtual HTTP—Virtual HTTP lets you authenticate separately with the security appliance and with the HTTP server. Even if the HTTP server does not need a second authentication, this feature achieves the effect of stripping the basic authentication credentials from the HTTP GET request. See the "Authenticating HTTP(S) Connections with a Virtual Server" section on page 24-7 for more information.
- Enabling the Exchange of Usernames and Passwords Using HTTPS—To enable the exchange of usernames and passwords between a web client and the security appliance with HTTPS, perform the following steps:

  a. From the Configuration > Firewall > AAA Rules pane, click **Advanced**. The AAA Rules Advanced Options dialog box appears.

  b. Under Secure HTTP, click **Enable Secure HTTP**.

  c. Click **OK**, and then click **OK** to exit the AAA Rules Advanced Options dialog box. Click **Apply**.

  This is the only method that protects credentials between the client and the security appliance, as well as between the security appliance and the destination server. You can use this method alone, or in conjunction with either of the other methods so you can maximize your security.

  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.

  Secured web-client authentication has the following limitations:

  – A maximum of 16 concurrent HTTPS authentication sessions are allowed. If all 16 HTTPS authentication processes are running, a new connection requiring authentication will not succeed.

  – When the uauth timeout is set to unlimited, HTTPS authentication might not work. If a browser initiates multiple TCP connections to load a web page after HTTPS authentication, the first connection is let through, but the subsequent connections trigger authentication. As a result, users are continuously presented with an authentication page, even if the correct username and password are entered each time. To work around this, set the uauth timeout to 1 second (see the Configuration > Firewall > Advanced > Global Timeouts pane). However, this workaround opens a 1-second window of opportunity that might allow non-authenticated users to go through the firewall if they are coming from the same source IP address.

  – Because HTTPS authentication occurs on the SSL port 443, users must not configure an Access Rule to block traffic from the HTTP client to HTTP server on port 443. Furthermore, if static PAT is configured for web traffic on port 80, it must also be configured for the SSL port.

Authenticating Directly with the Security Appliance

If you do not want to allow HTTP, HTTPS, Telnet, or FTP through the security appliance but want to authenticate other types of traffic, you can authenticate with the security appliance directly using HTTP, HTTPS, or Telnet.

- Authenticating Telnet Connections with a Virtual Server, page 24-7
- Authenticating HTTP(S) Connections with a Virtual Server, page 24-7

Authenticating Telnet Connections with a Virtual Server

Although you can configure network access authentication for any protocol or service (see the "Configuring Authentication for Network Access" section on page 24-1), you can authenticate directly with HTTP, Telnet, or FTP only. A user must first authenticate with one of these services before other traffic that requires authentication is allowed through. If you do not want to allow HTTP, Telnet, or FTP through the security appliance, but want to authenticate other types of traffic, you can configure virtual Telnet; the user Telnets to a given IP address configured on the security appliance, and the security appliance provides a Telnet prompt.

You must configure authentication for Telnet access to the virtual Telnet address as well as the other services you want to authenticate according to the "Configuring Authentication for Network Access" section on page 24-1.

When an unauthenticated user connects to the virtual Telnet IP address, the user is challenged for a username and password, and then authenticated by the AAA server. Once authenticated, the user sees the message "Authentication Successful." Then, the user can successfully access other services that require authentication.

For inbound users (from lower security to higher security), you must also include the virtual Telnet address as a destination interface in the Access Rule applied to the source interface. Moreover, you must add a static NAT rule for the virtual Telnet IP address, even if NAT is not required. An identity NAT rule is typically used (where you translate the address to itself).

For outbound users, there is an explicit permit for traffic, but if you apply an Access Rule to an inside interface, be sure to allow access to the virtual Telnet address. A static NAT rule is not required.

To logout from the security appliance, reconnect to the virtual Telnet IP address; you are prompted to log out.

To enable direct authentication using Telnet, perform the following steps:

Step 1 From the Configuration > Firewall > Advanced > Virtual Access > Virtual Telnet Server area, check the **Enable** check box.

Step 2 In the Virtual Telnet Server field, add the IP address of the virtual Telnet server.

Make sure this address is an unused address that is routed to the security appliance. For example, if you perform NAT for inside addresses accessing an outside server, and you want to provide outside access to the virtual HTTP server, you can use one of the global NAT addresses for the virtual HTTP server address.

Step 3 Click **Apply**.

The virtual server is added and the changes are saved to the running configuration.

Authenticating HTTP(S) Connections with a Virtual Server

When you use HTTP authentication on the security appliance (see the"Configuring Authentication for Network Access" section on page 24-1), the security appliance uses basic HTTP authentication by default. You can change the authentication method so that the security appliance redirects HTTP connections to web pages generated by the security appliance itself using the "Configuring HTTP Redirect" section on page 8-8.

However, if you continue to use basic HTTP authentication, then you might need the virtual HTTP server when you have cascading HTTP authentications.

If the destination HTTP server requires authentication in addition to the security appliance, then virtual HTTP lets you authenticate separately with the security appliance (via a AAA server) and with the HTTP server. Without virtual HTTP, the same username and password you used to authenticate with the security appliance is sent to the HTTP server; you are not prompted separately for the HTTP server username and password. Assuming the username and password is not the same for the AAA and HTTP servers, then the HTTP authentication fails.

This feature redirects all HTTP connections that require AAA authentication to the virtual HTTP server on the security appliance. The security appliance prompts for the AAA server username and password. After the AAA server authenticates the user, the security appliance redirects the HTTP connection back to the original server, but it does not include the AAA server username and password. Because the username and password are not included in the HTTP packet, the HTTP server prompts the user separately for the HTTP server username and password.

For inbound users (from lower security to higher security), you must also include the virtual HTTP address as a destination interface in the Access Rule applied to the source interface. Moreover, you must add a static NAT rule for the virtual HTTP IP address, even if NAT is not required. An identity NAT rule is typically used (where you translate the address to itself).

For outbound users, there is an explicit permit for traffic, but if you apply an Access Rule to an inside interface, be sure to allow access to the virtual HTTP address. A static NAT rule is not required.

Note Do not set the uauth timeout duration to 0 seconds when using virtual HTTP, because this setting prevents HTTP connections to the real web server. See the "Configuring Global Timeouts" section on page 28-26.

You can authenticate directly with the security appliance at the following URLs when you enable AAA for the interface:

```
http://interface_ip[:port]/netaccess/connstatus.html
https://interface_ip[:port]/netaccess/connstatus.html
```

To allow users to authenticate with the security appliance virtual server separately from the HTTP server, perform the following steps:

Step 1 From the Configuration > Firewall > Advanced > Virtual Access > Virtual HTTP Server area, check the **Enable** check box.

Step 2 In the Virtual HTTP Server field, add the IP address of the virtual HTTP server.

Make sure this address is an unused address that is routed to the security appliance. For example, if you perform NAT for inside addresses accessing an outside server, and you want to provide outside access to the virtual HTTP server, you can use one of the global NAT addresses for the virtual HTTP server address.

Step 3 (Optional) If you are using text-based browsers, where redirection does not happen automatically, check the **Display redirection warning** check box. This enables an alert to notify users when the HTTP connection is being redirected.

Step 4 Click **Apply**.

The virtual server is added and the changes are saved to the running configuration.

Configuring the Authentication Proxy Limit

You can manually configure the uauth session limit by setting the maximum number of concurrent proxy connections allowed per user.

To set the proxy limit, perform the following steps:

Step 1 From the Configuration > Firewall > AAA Rules pane, click **Advanced**.

The AAA Rules Advanced Options dialog box appears.

Step 2 In the Proxy Limit area, check **Enable Proxy Limit**.

Step 3 In the Proxy Limit field, enter the number of concurrent proxy connections allowed per user, from 1 to 128.

Step 4 Click **OK**, and then click **Apply**.

Configuring Authorization for Network Access

After a user authenticates for a given connection, the security appliance can use authorization to further control traffic from the user.

This section includes the following topics:

- Configuring TACACS+ Authorization, page 24-9
- Configuring RADIUS Authorization, page 24-10

Configuring TACACS+ Authorization

You can configure the security appliance to perform network access authorization with TACACS+.

Authentication and authorization rules are independent; however, any unauthenticated traffic matched by an authorization rule will be denied. For authorization to succeed:

1. 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 hasn’t expired, authorization can occur even if the traffic is not matched by an authentication rule.

2. After a user authenticates, the security appliance checks the authorization rules for matching traffic.
3. If the traffic matches the authorization rule, the security appliance sends the username to the TACACS+ server.
4. The TACACS+ server responds to the security appliance with a permit or a deny for that traffic, based on the user profile.
5. 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.

To configure TACACS+ authorization, perform the following steps:

Step 1 Enable authentication. For more information, see the "Configuring Network Access Authentication" section on page 24-4. If you have already enabled authentication, continue to the next step.

Step 2 From the Configuration > Firewall > AAA Rules pane, choose **Add > Add Authorization Rule**.

The Add Authorization Rule dialog box appears.

Step 3 From the Interface drop-down list, choose the interface for applying the rule.

Step 4 In the Action field, click one of the following, depending on the implementation:

- **Authorize**
- **Do not Authorize**.

Step 5 From the AAA Server Group drop-down list, choose a server group. To add a AAA server to the server group, click **Add Server**. See the "Configuring AAA Server Groups" section on page 16-9 for more information.

Only TACACS+ servers are supported.

Step 6 In the Source field, add the source IP address, or click the ellipsis (...) to choose an IP address already defined in ASDM.

Step 7 In the Destination field, enter the destination IP address, or click the ellipsis (...) to choose an IP address already defined in ASDM.

Step 8 In the Service field, enter an IP service name or number for the destination service, or click ellipsis (...) button to choose a service.

Step 9 (Optional) In the Description field, add a description.

Step 10 (Optional) Click **More Options** to do any of the following:

- To specify a source service for TCP or UDP, enter a TCP or UDP service in the Source Service field.

  The destination service and source service must be the same. Copy and paste the destination Service field to the Source Service field.
- To make the rule inactive, uncheck **Enable Rule**.

  You may not want to remove a rule, but instead turn it off.
- To set a time range for the rule, from the Time Range drop-down list, choose an existing time range. To add a new time range, click the ellipsis (...). For more information, see Configuring Time Ranges, page 20-15.

Step 11 Click **OK**.

The dialog box closes and the rule appears in the AAA Rules table.

Step 12 Click **Apply**.

The changes are saved to the running configuration.

Configuring RADIUS Authorization

When authentication succeeds, the RADIUS protocol returns user authorizations in the access-accept message sent by a RADIUS server. For more information about configuring authentication, see the "Configuring Authentication for Network Access" section on page 24-1.

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 enabled the Per User Override Setting (see the Configuration > Firewall > Access Rules > Advanced > Access Rules Advanced Options dialog box), be aware of the following effects of this feature on authorization by user-specific access lists:

- Without the per-user-override feature, 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 feature, the user-specific access list determines what is permitted.

This section includes the following topics:

- Configuring a RADIUS Server to Send Downloadable Access Control Lists, page 24-11
- Configuring a RADIUS Server to Download Per-User Access Control List Names, page 24-15

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, page 24-11
- Configuring Cisco Secure ACS for Downloadable Access Lists, page 24-13
- Configuring Any RADIUS Server for Downloadable Access Lists, page 24-14
- Converting Wildcard Netmask Expressions in Downloadable Access Lists, page 24-15

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 when you add a server to a server group, on the Configuration > Device Management > Users/AAA > AAA Server Groups > AAA Server Groups area. See the "Adding a Server to a Group" section on page 16-10.

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 (at the CLI) 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.

See the *Cisco ASA 5500 Series Configuration Guide using the CLI* to create an access list on the security appliance.

Configuring Accounting for Network Access

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.

To configure accounting, perform the following steps:

Step 1 If you want the security appliance to provide accounting data per user, you must enable authentication. For more information, see the "Configuring Network Access Authentication" section on page 24-4. If you want the security appliance to provide accounting data per IP address, enabling authentication is not necessary and you can continue to the next step.

Step 2 From the Configuration > Firewall > AAA Rules pane, choose **Add > Add Accounting Rule**.

The Add Accounting Rule dialog box appears.

Step 3 From the Interface drop-down list, choose the interface for applying the rule.

Step 4 In the Action field, click one of the following, depending on the implementation:

- **Account**
- **Do not Account**.

Step 5 From the AAA Server Group drop-down list, choose a server group. To add a AAA server to the server group, click **Add Server**. See the "Configuring AAA Server Groups" section on page 16-9 for more information.

Step 6 In the Source field, add the source IP address, or click the ellipsis (...) to choose an IP address already defined in ASDM.

Step 7 In the Destination field, enter the destination IP address, or click the ellipsis (...) to choose an IP address already defined in ASDM.

Step 8 In the Service field, enter an IP service name or number for the destination service, or click ellipsis (...) button to choose a service.

Step 9 (Optional) In the Description field, add a description.

Step 10 (Optional) Click **More Options** to do any of the following:

- To specify a source service for TCP or UDP, enter a TCP or UDP service in the Source Service field.

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

- To make the rule inactive, uncheck **Enable Rule**.

  You may not want to remove a rule, but instead turn it off.

- To set a time range for the rule, from the Time Range drop-down list, choose an existing time range. To add a new time range, click the ellipsis (...). For more information, see Configuring Time Ranges, page 20-15.

Step 11 Click **OK**.

The dialog box closes and the rule appears in the AAA Rules table.

Step 12 Click **Apply**.

The changes are saved to the running configuration.

Using MAC Addresses to Exempt Traffic from Authentication and Authorization

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. This feature is particularly useful to exempt devices such as IP phones that cannot respond to authentication prompts.

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.

To use MAC addresses to exempt traffic from authentication and authorization, perform the following steps:

Step 13 From the Configuration > Firewall > AAA Rules pane, choose **Add > Add MAC Exempt Rule**.

The Add MAC Exempt Rule dialog box appears.

Step 14 From the Action drop-down list, click one of the following, depending on the implementation:

- **MAC Exempt**
- **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.

Step 15 In the MAC Address field, specify the source MAC address in 12-digit hexadecimal form; that is, nnnn.nnnn.nnnn.

Step 16 In the MAC Mask field, specify 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.

Step 17 Click **OK**.

The dialog box closes and the rule appears in the AAA Rules table.

Step 18 Click **Apply**.

The changes are saved to the running configuration.

CHAPTER 25

Configuring Application Layer Protocol Inspection

This chapter describes how to configure application layer protocol inspection. Inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the security appliance to do a deep packet inspection instead of passing the packet through the fast path. As a result, inspection engines can affect overall throughput.

Several common inspection engines are enabled on the security appliance by default, but you might need to enable others depending on your network. This chapter includes the following sections:

Inspection Engine Overview

This section includes the following topics:

When to Use Application Protocol Inspection

When a user establishes a connection, the security appliance checks the packet against access lists, creates an address translation, and creates an entry for the session in the fast path, so that further packets can bypass time-consuming checks. However, the fast path relies on predictable port numbers and does not perform address translations inside a packet.

Many protocols open secondary TCP or UDP ports. The initial session on a well-known port is used to negotiate dynamically assigned port numbers.

Other applications embed an IP address in the packet that needs to match the source address that is normally translated when it goes through the security appliance.

If you use applications like these, then you need to enable application inspection.

When you enable application inspection for a service that embeds IP addresses, the security appliance translates embedded addresses and updates any checksum or other fields that are affected by the translation.

When you enable application inspection for a service that uses dynamically assigned ports, the security appliance monitors sessions to identify the dynamic port assignments, and permits data exchange on these ports for the duration of the specific session.

Inspection Limitations

See the following limitations for application protocol inspection:

- State information for multimedia sessions that require inspection are not passed over the state link for stateful failover. The exception is GTP, which is replicated over the state link.
- Some inspection engines do not support PAT, NAT, outside NAT, or NAT between same security interfaces. See "Default Inspection Policy" for more information about NAT support.

Default Inspection Policy

By default, the configuration includes a policy that matches all default application inspection traffic and applies inspection to the traffic on all interfaces (a global policy). Default application inspection traffic includes traffic to the default ports for each protocol. You can only apply one global policy, so if you want to alter the global policy, for example, to apply inspection to non-standard ports, or to add inspections that are not enabled by default, you need to either edit the default policy or disable it and apply a new one.

Table 25-1 lists all inspections supported, the default ports used in the default class map, and the inspection engines that are on by default, shown in bold. This table also notes any NAT limitations.

Table 25-1 Supported Application Inspection Engines

Application[1]	Default Port	NAT Limitations	Standards[2]	Comments
CTIQBE	TCP/2748	—	—	—
DCERPC	TCP/135	—	—	—
DNS over UDP	UDP/53	No NAT support is available for name resolution through WINS.	RFC 1123	No PTR records are changed.
FTP	TCP/21	—	RFC 959	—
GTP	UDP/3386 UDP/2123	—	—	Requires a special license.
H.323 H.225 and **RAS**	TCP/1720 UDP/1718 UDP (RAS) 1718-1719	No NAT on same security interfaces. No static PAT.	ITU-T H.323, H.245, H225.0, Q.931, Q.932	—
HTTP	TCP/80	—	RFC 2616	Beware of MTU limitations stripping ActiveX and Java. If the MTU is too small to allow the Java or ActiveX tag to be included in one packet, stripping may not occur.
ICMP	—	—	—	All ICMP traffic is matched in the default class map.
ICMP ERROR	—	—	—	All ICMP traffic is matched in the default class map.
ILS (LDAP)	TCP/389	No PAT.	—	—
Instant Messaging (IM)	Varies by client	—	RFC 3860	—

Table 25-1 Supported Application Inspection Engines (continued)

Application[1]	Default Port	NAT Limitations	Standards[2]	Comments
IP Options	—	—	RFC 791, RFC 2113	All IP Options traffic is matched in the default class map.
MMP	TCP 5443	—	—	—
MGCP	UDP/2427, 2727	—	RFC 2705bis-05	—
NetBIOS Name Server over IP	UDP/137, 138 (Source ports)	—	—	NetBIOS is supported by performing NAT of the packets for NBNS UDP port 137 and NBDS UDP port 138.
PPTP	TCP/1723	—	RFC 2637	—
RADIUS Accounting	1646	—	RFC 2865	—
RSH	TCP/514	No PAT	Berkeley UNIX	—
RTSP	TCP/554	No PAT. No outside NAT.	RFC 2326, 2327, 1889	No handling for HTTP cloaking.
SIP	TCP/5060 UDP/5060	No outside NAT. No NAT on same security interfaces.	RFC 2543	—
SKINNY (SCCP)	TCP/2000	No outside NAT. No NAT on same security interfaces.	—	Does not handle TFTP uploaded Cisco IP Phone configurations under certain circumstances.
SMTP and **ESMTP**	TCP/25	—	RFC 821, 1123	—
SNMP	UDP/161, 162	No NAT or PAT.	RFC 1155, 1157, 1212, 1213, 1215	v.2 RFC 1902-1908; v.3 RFC 2570-2580.
SQL*Net	TCP/1521	—	—	v.1 and v.2.
Sun RPC over UDP and TCP	UDP/111	No NAT or PAT.	—	The default rule includes UDP port 111; if you want to enable Sun RPC inspection for TCP port 111, you need to create a new rule that matches TCP port 111 and performs Sun RPC inspection.
TFTP	UDP/69	—	RFC 1350	Payload IP addresses are not translated.
WAAS	—	—	—	—
XDCMP	UDP/177	No NAT or PAT.	—	—

[1] Inspection engines that are enabled by default for the default port are in bold.

[2] The security appliance is in compliance with these standards, but it does not enforce compliance on packets being inspected. For example, FTP commands are supposed to be in a particular order, but the security appliance does not enforce the order.

Configuring Application Inspection

This feature uses Security Policy Rules. Service policies provide a consistent and flexible way to configure security appliance features. For example, you can use a service policy to create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications. See Chapter 23, "Configuring Service Policy Rules," for more information.

Inspection is enabled by default for some applications. See the "Default Inspection Policy" section for more information. Use this section to modify your inspection policy.

To configure application inspection, perform the following steps:

Step 1 Choose **Configuration > Firewall > Service Policy Rules**.

Step 2 Add or edit a service policy rule according to the "Adding a Service Policy Rule for Through Traffic" section on page 23-7.

If you want to match non-standard ports, then create a new rule for the non-standard ports. See the "Default Inspection Policy" section on page 25-3 for the standard ports for each inspection engine. You can combine multiple rules in the same service policy if desired, so you can create one rule to match certain traffic, and another to match different traffic. However, if traffic matches a rule that contains an inspection action, and then matches another rule that also has an inspection action, only the first matching rule is used.

Step 3 In the Edit Service Policy Rule > Rule Actions dialog box, click the **Protocol Inspection** tab.

For a new rule, the dialog box is called Add Service Policy Rule Wizard - Rule Actions.

Step 4 Select each inspection type that you want to apply.

Step 5 (Optional) Some inspection engines let you control additional parameters when you apply the inspection to the traffic. Click **Configure** for each inspection type to configure an inspect map.

You can either choose an existing map, or create a new one. You can predefine inspect maps in the Configuration > Firewall > Objects > Inspect Maps pane. See the "Inspect Map Field Descriptions" section on page 25-64 for detailed information of each inspect map type.

Step 6 You can configure other features for this rule if desired using the other Rule Actions tabs.

Step 7 Click **OK** (or **Finish** from the wizard).

CTIQBE Inspection

This section describes CTIQBE application inspection. This section includes the following topics:

- CTIQBE Inspection Overview, page 25-5
- Limitations and Restrictions, page 25-6

CTIQBE Inspection Overview

CTIQBE protocol inspection supports NAT, PAT, and bidirectional NAT. This enables Cisco IP SoftPhone and other Cisco TAPI/JTAPI applications to work successfully with Cisco CallManager for call setup across the security appliance.

TAPI and JTAPI are used by many Cisco VoIP applications. CTIQBE is used by Cisco TSP to communicate with Cisco CallManager.

Limitations and Restrictions

The following summarizes limitations that apply when using CTIQBE application inspection:

- CTIQBE application inspection does not support configurations with the **alias** command.
- Stateful failover of CTIQBE calls is not supported.
- Entering the **debug ctiqbe** command may delay message transmission, which may have a performance impact in a real-time environment. When you enable this debugging or logging and Cisco IP SoftPhone seems unable to complete call setup through the security appliance, increase the timeout values in the Cisco TSP settings on the system running Cisco IP SoftPhone.

The following summarizes special considerations when using CTIQBE application inspection in specific scenarios:

- If two Cisco IP SoftPhones are registered with different Cisco CallManagers, which are connected to different interfaces of the security appliance, calls between these two phones fails.
- When Cisco CallManager is located on the higher security interface compared to Cisco IP SoftPhones, if NAT or outside NAT is required for the Cisco CallManager IP address, the mapping must be static as Cisco IP SoftPhone requires the Cisco CallManager IP address to be specified explicitly in its Cisco TSP configuration on the PC.
- When using PAT or Outside PAT, if the Cisco CallManager IP address is to be translated, its TCP port 2748 must be statically mapped to the same port of the PAT (interface) address for Cisco IP SoftPhone registrations to succeed. The CTIQBE listening port (TCP 2748) is fixed and is not user-configurable on Cisco CallManager, Cisco IP SoftPhone, or Cisco TSP.

DCERPC 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 listening on a well known port number for the dynamically allocated network information of a required service. The client then sets up a secondary connection to the server instance providing the service. The security appliance allows the appropriate port number and network address and also applies NAT, if needed, for the secondary connection.

DCERPC inspect maps inspect for native TCP communication between the EPM and client on well known TCP port 135. Map and lookup operations of the EPM are supported for clients. Client and server can be located in any security zone. The embedded server IP address and Port number are received from the applicable EPM response messages. Because a client may attempt multiple connections to the server port returned by EPM, multiple use of pinholes are allowed, which have user configurable timeouts.

DNS Inspection

This section describes DNS application inspection. This section includes the following topics:

- How DNS Application Inspection Works, page 25-7

- How DNS Rewrite Works, page 25-7

How DNS Application Inspection Works

The security appliance tears down the DNS session associated with a DNS query as soon as the DNS reply is forwarded by the security appliance. The security appliance also monitors the message exchange to ensure that the ID of the DNS reply matches the ID of the DNS query.

When DNS inspection is enabled, which is the default, the security appliance performs the following additional tasks:

- Translates the DNS record based on the configuration completed using NAT rules. Translation only applies to the A-record in the DNS reply; therefore, DNS Rewrite does not affect reverse lookups, which request the PTR record.

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

- Enforces the maximum DNS message length (the default is 512 bytes and the maximum length is 65535 bytes). The security appliance performs reassembly as needed to verify that the packet length is less than the maximum length configured. The security appliance drops the packet if it exceeds the maximum length.
- Enforces a domain-name length of 255 bytes and a label length of 63 bytes.
- Verifies the integrity of the domain-name referred to by the pointer if compression pointers are encountered in the DNS message.
- Checks to see if a compression pointer loop exists.

A single connection is created for multiple DNS sessions, as long as they are between the same two hosts, and the sessions have the same 5-tuple (source/destination IP address, source/destination port, and protocol). DNS identification is tracked by *app_id*, and the idle timer for each app_id runs independently.

Because the app_id expires independently, a legitimate DNS response can only pass through the security appliance within a limited period of time and there is no resource build-up. However, if you enter the **show conn** command, you will see the idle timer of a DNS connection being reset by a new DNS session. This is due to the nature of the shared DNS connection and is by design.

How DNS Rewrite Works

When DNS inspection is enabled, DNS rewrite provides full support for NAT of DNS messages originating from any interface.

If a client on an inside network requests DNS resolution of an inside address from a DNS server on an outside interface, the DNS A-record is translated correctly. If the DNS inspection engine is disabled, the A-record is not translated.

As long as DNS inspection remains enabled, you can configure DNS rewrite using a NAT rule.

DNS Rewrite performs two functions:

- Translating a public address (the routable or “mapped” address) in a DNS reply to a private address (the “real” address) when the DNS client is on a private interface.

- Translating a private address to a public address when the DNS client is on the public interface.

In Figure 25-1, the DNS server resides on the external (ISP) network The real address of the server (192.168.100.1) has been mapped using a static NAT rule to the ISP-assigned address (209.165.200.5). When a web client on the inside interface attempts to access the web server with the URL http://server.example.com, the host running the web client sends a DNS request to the DNS server to resolve the IP address of the web server. The security appliance translates the non-routable source address in the IP header and forwards the request to the ISP network on its outside interface. When the DNS reply is returned, the security appliance applies address translation not only to the destination address, but also to the embedded IP address of the web server, which is contained in the A-record in the DNS reply. As a result, the web client on the inside network gets the correct address for connecting to the web server on the inside network.

Figure 25-1 Translating the Address in a DNS Reply (DNS Rewrite)

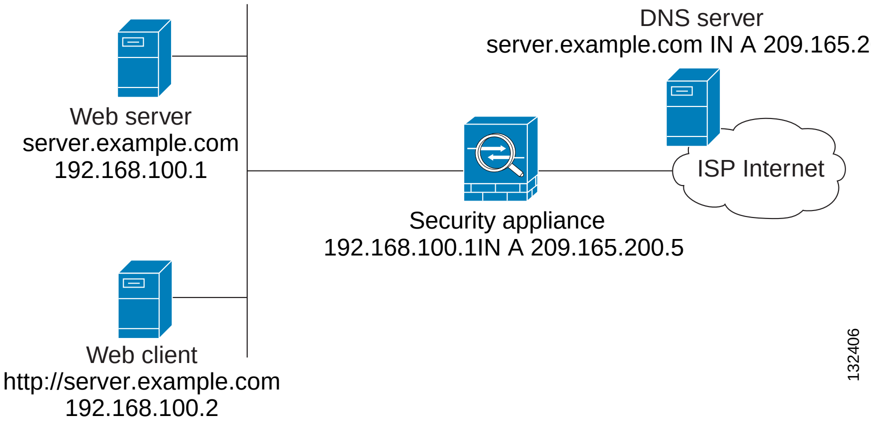

DNS rewrite also works if the client making the DNS request is on a DMZ network and the DNS server is on an inside interface.

ESMTP Inspection

ESMTP inspection detects attacks, including spam, phising, malformed message attacks, buffer overflow/underflow attacks. It also provides support for application security and protocol conformance, which enforce the sanity of the ESMTP messages as well as detect several attacks, block senders/receivers, and block mail relay.

FTP Inspection

This section describes the FTP inspection engine. This section includes the following topics:

- FTP Inspection Overview, page 25-9
- Using Strict FTP, page 25-9
- Verifying and Monitoring FTP Inspection, page 25-10

FTP Inspection Overview

The FTP application inspection inspects the FTP sessions and performs four tasks:

- Prepares dynamic secondary data connection
- Tracks the FTP command-response sequence
- Generates an audit trail
- Translates the embedded IP address

FTP application inspection prepares secondary channels for FTP data transfer. Ports for these channels are negotiated through PORT or PASV commands. The channels are allocated in response to a file upload, a file download, or a directory listing event.

Note If you disable FTP inspection engines, outbound users can start connections only in passive mode, and all inbound FTP is disabled.

Using Strict FTP

Using strict FTP increases the security of protected networks by preventing web browsers from sending embedded commands in FTP requests. To enable strict FTP, click the **Configure** button next to FTP on the Configuration > Firewall > Service Policy Rules > Edit Service Policy Rule > Rule Actions > Protocol Inspection tab.

Note To specify FTP commands that are not permitted to pass through the security appliance, create an FTP inspect map according to the "FTP Class Map" section on page 25-47.

After you enable the Strict option on an interface, FTP inspection enforces the following behavior:

- An FTP command must be acknowledged before the security appliance allows a new command.
- The security appliance drops connections that send embedded commands.
- The 227 and PORT commands are checked to ensure they do not appear in an error string.

Caution Using the strict option may cause the failure of FTP clients that are not strictly compliant with FTP RFCs.

If the strict option is enabled, each FTP command and response sequence is tracked for the following anomalous activity:

- Truncated command—Number of commas in the PORT and PASV reply command is checked to see if it is five. If it is not five, then the PORT command is assumed to be truncated and the TCP connection is closed.
- Incorrect command—Checks the FTP command to see if it ends with <CR><LF> characters, as required by the RFC. If it does not, the connection is closed.
- Size of RETR and STOR commands—These are checked against a fixed constant. If the size is greater, then an error message is logged and the connection is closed.
- Command spoofing—The PORT command should always be sent from the client. The TCP connection is denied if a PORT command is sent from the server.

- Reply spoofing—PASV reply command (227) should always be sent from the server. The TCP connection is denied if a PASV reply command is sent from the client. This prevents the security hole when the user executes “227 xxxxx a1, a2, a3, a4, p1, p2.”
- TCP stream editing—The security appliance closes the connection if it detects TCP stream editing.
- Invalid port negotiation—The negotiated dynamic port value is checked to see if it is less than 1024. As port numbers in the range from 1 to 1024 are reserved for well-known connections, if the negotiated port falls in this range, then the TCP connection is freed.
- Command pipelining—The number of characters present after the port numbers in the PORT and PASV reply command is cross checked with a constant value of 8. If it is more than 8, then the TCP connection is closed.
- The security appliance replaces the FTP server response to the SYST command with a series of Xs. to prevent the server from revealing its system type to FTP clients. To override this default behavior, use the Low setting in the FTP map.

Verifying and Monitoring FTP Inspection

FTP application inspection generates the following log messages:

- An Audit record 302002 is generated for each file that is retrieved or uploaded.
- The FTP command is checked to see if it is RETR or STOR and the retrieve and store commands are logged.
- The username is obtained by looking up a table providing the IP address.
- The username, source IP address, destination IP address, NAT address, and the file operation are logged.
- Audit record 201005 is generated if the secondary dynamic channel preparation failed due to memory shortage.

In conjunction with NAT, the FTP application inspection translates the IP address within the application payload. This is described in detail in RFC 959.

GTP Inspection

Note GTP inspection requires a special license.

GPRS provides uninterrupted connectivity for mobile subscribers between GSM networks and corporate networks or the Internet. The GGSN is the interface between the GPRS wireless data network and other networks. The SGSN performs mobility, data session management, and data compression (See Figure 25-2).

Figure 25-2 GPRS Tunneling Protocol

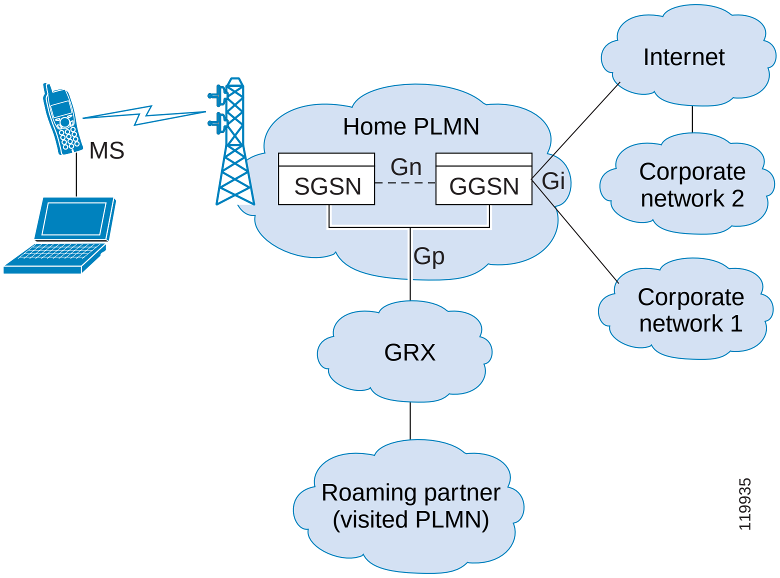

The UMTS is the commercial convergence of fixed-line telephony, mobile, Internet and computer technology. UTRAN is the networking protocol used for implementing wireless networks in this system. GTP allows multi-protocol packets to be tunneled through a UMTS/GPRS backbone between a GGSN, an SGSN and the UTRAN.

GTP does not include any inherent security or encryption of user data, but using GTP with the security appliance helps protect your network against these risks.

The SGSN is logically connected to a GGSN using GTP. GTP allows multiprotocol packets to be tunneled through the GPRS backbone between GSNs. GTP provides a tunnel control and management protocol that allows the SGSN to provide GPRS network access for a mobile station by creating, modifying, and deleting tunnels. GTP uses a tunneling mechanism to provide a service for carrying user data packets.

Note When using GTP with failover, if a GTP connection is established and the active unit fails before data is transmitted over the tunnel, the GTP data connection (with a “j” flag set) is not replicated to the standby unit. This occurs because the active unit does not replicate embryonic connections to the standby unit.

H.323 Inspection

This section describes the H.323 application inspection. This section includes the following topics:

- H.323 Inspection Overview, page 25-12
- How H.323 Works, page 25-12
- H.239 Support in H.245 Messages, page 25-13
- Limitations and Restrictions, page 25-13

H.323 Inspection Overview

H.323 inspection provides support for H.323 compliant applications such as Cisco CallManager and VocalTec Gatekeeper. H.323 is a suite of protocols defined by the International Telecommunication Union for multimedia conferences over LANs. The security appliance supports H.323 through Version 6, including H.323 v3 feature Multiple Calls on One Call Signaling Channel.

With H.323 inspection enabled, the security appliance supports multiple calls on the same call signaling channel, a feature introduced with H.323 Version 3. This feature reduces call setup time and reduces the use of ports on the security appliance.

The two major functions of H.323 inspection are as follows:

- NAT the necessary embedded IPv4 addresses in the H.225 and H.245 messages. Because H.323 messages are encoded in PER encoding format, the security appliance uses an ASN.1 decoder to decode the H.323 messages.
- Dynamically allocate the negotiated H.245 and RTP/RTCP connections.

How H.323 Works

The H.323 collection of protocols collectively may use up to two TCP connection and four to eight UDP connections. FastConnect uses only one TCP connection, and RAS uses a single UDP connection for registration, admissions, and status.

An H.323 client may initially establish a TCP connection to an H.323 server using TCP port 1720 to request Q.931 call setup. As part of the call setup process, the H.323 terminal supplies a port number to the client to use for an H.245 TCP connection. In environments where H.323 gatekeeper is in use, the initial packet is transmitted using UDP.

H.323 inspection monitors the Q.931 TCP connection to determine the H.245 port number. If the H.323 terminals are not using FastConnect, the security appliance dynamically allocates the H.245 connection based on the inspection of the H.225 messages.

Note The H.225 connection can also be dynamically allocated when using RAS.

Within each H.245 message, the H.323 endpoints exchange port numbers that are used for subsequent UDP data streams. H.323 inspection inspects the H.245 messages to identify these ports and dynamically creates connections for the media exchange. RTP uses the negotiated port number, while RTCP uses the next higher port number.

The H.323 control channel handles H.225 and H.245 and H.323 RAS. H.323 inspection uses the following ports.

- 1718—Gate Keeper Discovery UDP port
- 1719—RAS UDP port
- 1720—TCP Control Port

You must permit traffic for the well-know H.323 port 1719 for RAS signaling. Additionally, you must permit traffic for the well-known H.323 port 1720 for the H.225 call signaling; however, the H.245 signaling ports are negotiated between the endpoints in the H.225 signaling. When an H.323 gatekeeper is used, the security appliance opens an H.225 connection based on inspection of the ACF and RCF messages.

Note You can enable call setup between H.323 endpoints when the Gatekeeper is inside the network. The security appliance includes options to open pinholes for calls based on the RegistrationRequest/RegistrationConfirm (RRQ/RCF) messages. Because these RRQ/RCF messages are sent to and from the Gatekeeper, the calling endpoint's IP address is unknown and the security appliance opens a pinhole through source IP address/port 0/0. By default, this option is disabled. You can enable this option by setting the option in the H.323 Inspect Map. From ASDM, choose Configuration > Firewall > Objects > Inspect Maps > H.323 > Details > State Checking tab.

After inspecting the H.225 messages, the security appliance opens the H.245 channel and then inspects traffic sent over the H.245 channel as well. All H.245 messages passing through the security appliance undergo H.245 application inspection, which translates embedded IP addresses and opens the media channels negotiated in H.245 messages.

The H.323 ITU standard requires that a TPKT header, defining the length of the message, precede the H.225 and H.245, before being passed on to the reliable connection. Because the TPKT header does not necessarily need to be sent in the same TCP packet as H.225 and H.245 messages, the security appliance must remember the TPKT length to process and decode the messages properly. For each connection, the security appliance keeps a record that contains the TPKT length for the next expected message.

If the security appliance needs to perform NAT on IP addresses in messages, it changes the checksum, the UUIE length, and the TPKT, if it is included in the TCP packet with the H.225 message. If the TPKT is sent in a separate TCP packet, the security appliance proxy ACKs that TPKT and appends a new TPKT to the H.245 message with the new length.

Note The security appliance does not support TCP options in the Proxy ACK for the TPKT.

Each UDP connection with a packet going through H.323 inspection is marked as an H.323 connection and times out with the H.323 timeout as configured in the Configuration > Firewall > Advanced > Global Timeouts pane.

H.239 Support in H.245 Messages

The security appliance sits between two H.323 endpoints. When the two H.323 endpoints set up a telepresentation session so that the endpoints can send and receive a data presentation, such as spreadsheet data, the security appliance ensure successful H.239 negotiation between the endpoints.

H.239 is a standard that provides the ability for H.300 series endpoints to open an additional video channel in a single call. In a call, an endpoint (such as a video phone), sends a channel for video and a channel for data presentation. The H.239 negotiation occurs on the H.245 channel.

The security appliance opens pinholes for the additional media channel and the media control channel. The endpoints use open logical channel message (OLC) to signal a new channel creation. The message extension is part of H.245 Version 13.

The decoding and encoding of of the telepresentation session is enabled by default. H.239 encoding and decoding is preformed by ASN.1 coder.

Limitations and Restrictions

The following are some of the known issues and limitations when using H.323 application inspection:

- Static PAT may not properly translate IP addresses embedded in optional fields within H.323 messages. If you experience this kind of problem, do not use static PAT with H.323.
- H.323 application inspection is not supported with NAT between same-security-level interfaces.
- When a NetMeeting client registers with an H.323 gatekeeper and tries to call an H.323 gateway that is also registered with the H.323 gatekeeper, the connection is established but no voice is heard in either direction. This problem is unrelated to the security appliance.
- If you configure a network static address where the network static address is the same as a third-party netmask and address, then any outbound H.323 connection fails.

HTTP Inspection

Use the HTTP inspection engine to protect against specific attacks and other threats that may be associated with HTTP traffic. HTTP inspection performs several functions:

- Enhanced HTTP inspection
- URL screening through N2H2 or Websense
- Java and ActiveX filtering

The latter two features are configured in conjunction with Filter rules.

The enhanced HTTP inspection feature, which is also known as an application firewall and is available when you configure an HTTP inspect map (see the "HTTP Class Map" section on page 25-52), can help prevent attackers from using HTTP messages for circumventing network security policy. It verifies the following for all HTTP messages:

- Conformance to RFC 2616
- Use of RFC-defined methods only.
- Compliance with the additional criteria.

Instant Messaging Inspection

The IM inspect engine lets you apply fine grained controls on the IM application to control the network usage and stop leakage of confidential data, propagation of worms, and other threats to the corporate network.

ICMP Inspection

The ICMP inspection engine allows ICMP traffic to have a "session" so it can be inspected like TCP and UDP traffic. Without the ICMP inspection engine, we recommend that you do not allow ICMP through the security appliance in an access list. Without stateful inspection, ICMP can be used to attack your network. The ICMP inspection engine ensures that there is only one response for each request, and that the sequence number is correct.

ICMP Error Inspection

When this feature is enabled, the security appliance creates translation sessions for intermediate hops that send ICMP error messages, based on the NAT configuration. The security appliance overwrites the packet with the translated IP addresses.

When disabled, the security appliance does not create translation sessions for intermediate nodes that generate ICMP error messages. ICMP error messages generated by the intermediate nodes between the inside host and the security appliance reach the outside host without consuming any additional NAT resource. This is undesirable when an outside host uses the traceroute command to trace the hops to the destination on the inside of the security appliance. When the security appliance does not translate the intermediate hops, all the intermediate hops appear with the mapped destination IP address.

The ICMP payload is scanned to retrieve the five-tuple from the original packet. Using the retrieved five-tuple, a lookup is performed to determine the original address of the client. The ICMP error inspection engine makes the following changes to the ICMP packet:

- In the IP Header, the mapped IP is changed to the real IP (Destination Address) and the IP checksum is modified.
- In the ICMP Header, the ICMP checksum is modified due to the changes in the ICMP packet.
- In the Payload, the following changes are made:
  - Original packet mapped IP is changed to the real IP
  - Original packet mapped port is changed to the real Port
  - Original packet IP checksum is recalculated

ILS Inspection

The ILS inspection engine provides NAT support for Microsoft NetMeeting, SiteServer, and Active Directory products that use LDAP to exchange directory information with an ILS server.

The security appliance supports NAT for ILS, which is used to register and locate endpoints in the ILS or SiteServer Directory. PAT cannot be supported because only IP addresses are stored by an LDAP database.

For search responses, when the LDAP server is located outside, NAT should be considered to allow internal peers to communicate locally while registered to external LDAP servers. For such search responses, xlates are searched first, and then DNAT entries to obtain the correct address. If both of these searches fail, then the address is not changed. For sites using NAT 0 (no NAT) and not expecting DNAT interaction, we recommend that the inspection engine be turned off to provide better performance.

Additional configuration may be necessary when the ILS server is located inside the security appliance border. This would require a hole for outside clients to access the LDAP server on the specified port, typically TCP 389.

Because ILS traffic only occurs on the secondary UDP channel, the TCP connection is disconnected after the TCP inactivity interval. By default, this interval is 60 minutes and can be adjusted using the Configuration > Firewall > Advanced > Global Timeouts pane.

ILS/LDAP follows a client/server model with sessions handled over a single TCP connection. Depending on the client's actions, several of these sessions may be created.

During connection negotiation time, a BIND PDU is sent from the client to the server. Once a successful BIND RESPONSE from the server is received, other operational messages may be exchanged (such as ADD, DEL, SEARCH, or MODIFY) to perform operations on the ILS Directory. The ADD REQUEST and SEARCH RESPONSE PDUs may contain IP addresses of NetMeeting peers, used by H.323 (SETUP and CONNECT messages) to establish the NetMeeting sessions. Microsoft NetMeeting v2.X and v3.X provides ILS support.

The ILS inspection performs the following operations:

- Decodes the LDAP REQUEST/RESPONSE PDUs using the BER decode functions
- Parses the LDAP packet
- Extracts IP addresses
- Translates IP addresses as necessary
- Encodes the PDU with translated addresses using BER encode functions
- Copies the newly encoded PDU back to the TCP packet
- Performs incremental TCP checksum and sequence number adjustment

ILS inspection has the following limitations:

- Referral requests and responses are not supported
- Users in multiple directories are not unified
- Single users having multiple identities in multiple directories cannot be recognized by NAT

Note Because H.225 call signalling traffic only occurs on the secondary UDP channel, the TCP connection is disconnected after the interval specified by the TCP option in the Configuration > Firewall > Advanced > Global Timeouts pane. By default, this interval is set at 60 minutes.

IP Options Inspection

In a packet, the IP header contains the Options field. The Options field, commonly referred to as IP Options, provide for control functions that are required in some situations but unnecessary for most common communications. In particular, IP Options include provisions for time stamps, security, and special routing. Use of IP Options is optional and the field can contain zero, one, or more options.

You can configure IP Options inspection to control which IP packets with specific IP options are allowed through the security appliance. Configuring this inspection instructs the security appliance to allow a packet to pass or to clear the specified IP options and then allow the packet to pass.

IP Options inspection can check for the following three IP options in a packet:

- End of Options List (EOOL) or IP Option 0—This option, which contains just a single zero byte, appears at the end of all options to mark the end of a list of options. This might not coincide with the end of the header according to the header length.
- No Operation (NOP) or IP Option 1—The Options field in the IP header can contain zero, one, or more options, which makes the total length of the field variable. However, the IP header must be a multiple of 32 bits. If the number of bits of all options is not a multiple of 32 bits, the NOP option is used as “internal padding” to align the options on a 32-bit boundary.
- Router Alert (RTRALT) or IP Option 20—This option notifies transit routers to inspect the contents of the packet even when the packet is not destined for that router. This inspection is valuable when implementing RSVP and similar protocols require relatively complex processing from the routers along the packets delivery path.

Note IP Options inspection is included by default in the global inspection policy. Therefore, the security appliance allows RSVP traffic that contains packets with the Router Alert option (option 20) when the security appliance is in routed mode.

Dropping RSVP packets containing the Router Alert option can cause problems in VoIP implementations.

When you configure security appliance to clear the Router Alert option from IP headers, the IP header changes in the following ways:

- The Options field is padded so that the field ends on a 32 bit boundary.
- Internet header length (IHL) changes.
- The total length of the packet changes.
- The checksum is recomputed.

If an IP header contains additional options other than EOOL, NOP, or RTRALT, regardless of whether the security appliance is configured to allow these options, the security appliance will drop the packet.

Configuring IP Options Inspection

Use the Add Service Policy Rule Wizard - Rule Actions dialog box to configure IP Options inspection.

This wizard is available from the Configuration > Firewall > Service Policy Rules > Add > Add Service Policy Rule Wizard - Rule Actions dialog box.

Step 1 Open the Add Service Policy Rule Wizard by selecting Configuration > Firewall > Service Policy Rules > Add. Perform the steps to complete the Service Policy, Traffic Classification Criteria, and Traffic Match - Destination Port pages of the wizard. See Adding a Service Policy Rule for Through Traffic, page 23-7.

The Add Service Policy Rule Wizard - Rule Actions dialog box opens.

Step 2 Check the IP-Options check box.

Step 3 Click **Configure** beside to the IP-Options check box. The Select IP Options Inspect Map dialog box opens.

Step 4 Perform one of the following:

- Click the **Use the default IP-Options inspection map** radio button to use the default IP Options map. The default map drops packets containing all the inspected IP options, namely End of Options List (EOOL), No Operation (NOP), and Router Alert (RTRALT).
- Click the **Select an IP-Options inspect map for fine control over inspection** radio button to select a defined application inspection map.
- Click Add to open the Add IP-Options Inspect Map dialog box and create a new inspection map.

Step 5 (Optional) If you clicked **Add** to create a new inspection map, define the following values for IP Options Inspection:

a. Enter a name for the inspection map.

b. Enter a description for the inspection map, up to 200 characters long.

c. From the Parameters area, select which IP options you want to pass through the security appliance or clear and then pass through the security appliance:

- Allow packets with the End of Options List (EOOL) option

  This option, which contains just a single zero byte, appears at the end of all options to mark the end of a list of options. This might not coincide with the end of the header according to the header length.

- Allow packets with the No Operation (NOP) option

  The Options field in the IP header can contain zero, one, or more options, which makes the total length of the field variable. However, the IP header must be a multiple of 32 bits. If the number of bits of all options is not a multiple of 32 bits, the NOP option is used as “internal padding” to align the options on a 32-bit boundary.

- Allow packets with the Router Alert (RTRALT) option

  This option notifies transit routers to inspect the contents of the packet even when the packet is not destined for that router. This inspection is valuable when implementing RSVP and similar protocols require relatively complex processing from the routers along the packets delivery path.

- Clear the option value from the packets

  When an option is checked, the **Clear the option value from the packets** check box becomes available for that option. Select the **Clear the option value from the packets** check box to clear the option from the packet before allowing the packet through the security appliance.

d. Click **OK**.

Step 6 Click **OK**.

Step 7 Click **Finish**.

MGCP Inspection

MGCP is a master/slave protocol used to control media gateways from external call control elements called media gateway controllers or call agents. A media gateway is typically a network element that provides conversion between the audio signals carried on telephone circuits and data packets carried over the Internet or over other packet networks. Using NAT and PAT with MGCP lets you support a large number of devices on an internal network with a limited set of external (global) addresses. Examples of media gateways are:

- Trunking gateways, that interface between the telephone network and a Voice over IP network. Such gateways typically manage a large number of digital circuits.
- Residential gateways, that provide a traditional analog (RJ11) interface to a Voice over IP network. Examples of residential gateways include cable modem/cable set-top boxes, xDSL devices, broad-band wireless devices.
- Business gateways, that provide a traditional digital PBX interface or an integrated soft PBX interface to a Voice over IP network.

MGCP messages are transmitted over UDP. A response is sent back to the source address (IP address and UDP port number) of the command, but the response may not arrive from the same address as the command was sent to. This can happen when multiple call agents are being used in a failover configuration and the call agent that received the command has passed control to a backup call agent, which then sends the response. Figure 25-3 illustrates how NAT can be used with MGCP.

Figure 25-3 Using NAT with MGCP

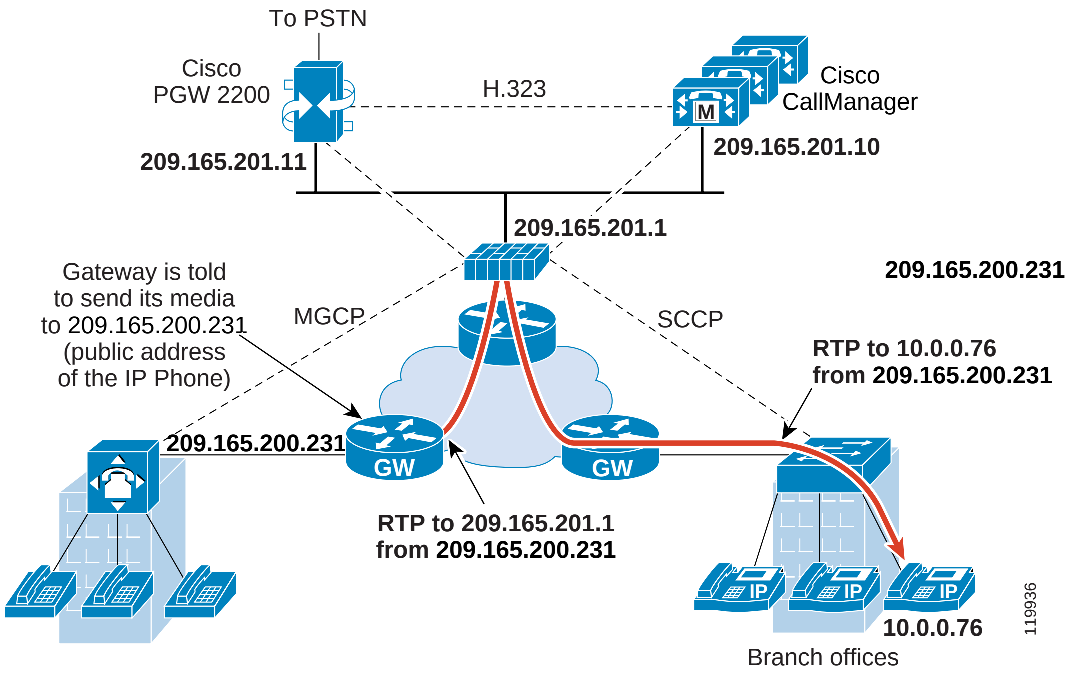
To PSTN
Cisco PGW 2200
H.323
Cisco CallManager
M
209.165.201.11
209.165.201.10
209.165.201.1
209.165.200.231
Gateway is told to send its media to 209.165.200.231 (public address of the IP Phone)
MGCP
SCCP
RTP to 10.0.0.76 from 209.165.200.231
209.165.200.231
GW
GW
RTP to 209.165.201.1 from 209.165.200.231
IP
IP
IP
10.0.0.76
Branch offices
119936

MGCP endpoints are physical or virtual sources and destinations for data. Media gateways contain endpoints on which the call agent can create, modify and delete connections to establish and control media sessions with other multimedia endpoints. Also, the call agent can instruct the endpoints to detect certain events and generate signals. The endpoints automatically communicate changes in service state to the call agent.

MGCP transactions are composed of a command and a mandatory response. There are eight types of commands:

- CreateConnection
- ModifyConnection
- DeleteConnection
- NotificationRequest
- Notify
- AuditEndpoint
- AuditConnection
- RestartInProgress

The first four commands are sent by the call agent to the gateway. The Notify command is sent by the gateway to the call agent. The gateway may also send a DeleteConnection. The registration of the MGCP gateway with the call agent is achieved by the RestartInProgress command. The AuditEndpoint and the AuditConnection commands are sent by the call agent to the gateway.

All commands are composed of a Command header, optionally followed by a session description. All responses are composed of a Response header, optionally followed by a session description.

- The port on which the gateway receives commands from the call agent. Gateways usually listen to UDP port 2427.

- The port on which the call agent receives commands from the gateway. Call agents usually listen to UDP port 2727.

Note MGCP inspection does not support the use of different IP addresses for MGCP signaling and RTP data. A common and recommended practice is to send RTP data from a resilient IP address, such as a loopback or virtual IP address; however, the security appliance requires the RTP data to come from the same address as MGCP signalling.

MMP Inspection

The security appliance includes an inspection engine to validate the CUMA Mobile Multiplexing Protocol (MMP).

For information about setting up the TLS Proxy for the Mobility Advantage feature, seethe "TLS Proxy Wizard" section on page 25-133.

MMP is a data transport protocol for transmitting data entities between CUMA clients and servers. MMP must be run on top of a connection-oriented protocol (the underlying transport) and is intended to be run on top of a secure transport protocol such as TLS. The Orative Markup Language (OML) protocol is intended to be run on top of MMP for the purposes of data synchronization, as well as the HTTP protocol for uploading and downloading large files.

The TCP/TLS default port is 5443. There are no embedded NAT or secondary connections.

CUMA client and server communications can be proxied via TLS, which decrypts the data, passes it to the inspect MMP module, and re-encrypt the data before forwarding it to the endpoint. The inspect MMP module verifies the integrity of the MMP headers and passes the OML/HTTP to an appropriate handler. The security appliance takes the following actions on the MMP headers and data:

- Verifies that client MMP headers are well-formed. Upon detection of a malformed header, the TCP session is terminated.
- Verifies that client to server MMP header lengths are not exceeded. If an MMP header length is exceeded (4096), then the TCP session is terminated.
- Verifies that client to server MMP content lengths are not exceeded. If an entity content length is exceeded (4096), the TCP session is terminated.

Note 4096 is the value currently used in MMP implementations.

Because MMP headers and entities can be split across packets, the security appliance buffers data to ensure consistent inspection. The SAPI (stream API) handles data buffering for pending inspection opportunities. MMP header text is treated as case insensitive and a space is present between header text and values. Reclaiming of MMP state is performed by monitoring the state of the TCP connection. Timeouts for these connections follow existing configurable values via the **timeout** command.

MMP inspection is disabled by default. When enabled, MMP inspection operates on TCP destination and source port 5443.

Configuring MMP Inspection for a TLS Proxy

Use the Add Service Policy Rule Wizard - Rule Actions dialog box to configure MMP protocol inspection.

This wizard is available from the Configuration > Firewall > Service Policy Rules > Add > Add Service Policy Rule Wizard - Rule Actions dialog box.

Step 1 Open the Add Service Policy Rule Wizard by selecting Configuration > Firewall > Service Policy Rules > Add. Perform the steps to complete the Service Policy, Traffic Classification Criteria, and Traffic Match - Destination Port pages of the wizard. See Adding a Service Policy Rule for Through Traffic, page 23-7.

The Add Service Policy Rule Wizard - Rule Actions dialog box opens.

Step 2 Check the MMP check box.

Step 3 Click **Configure** beside to the MMP check box. The Configure TLS Proxy dialog box opens.

Step 4 Perform one of the following:

Select the TLS Proxy for which you are enabling MMP protocol inspection.

Or

Click **Manage** to create a new TLS Proxy Instance. The Configure TLS Proxy dialog box opens. See the “Configure TLS Proxy Pane” section on page 25-134.

Step 5 Click **OK**.

Step 6 Click **Finish**.

NetBIOS Inspection

NetBIOS inspection is enabled by default. The NetBIOS inspection engine translates IP addresses in the NetBIOS name service (NBNS) packets according to the security appliance NAT configuration.

PPTP Inspection

PPTP is a protocol for tunneling PPP traffic. A PPTP session is composed of one TCP channel and usually two PPTP GRE tunnels. The TCP channel is the control channel used for negotiating and managing the PPTP GRE tunnels. The GRE tunnels carries PPP sessions between the two hosts.

When enabled, PPTP application inspection inspects PPTP protocol packets and dynamically creates the GRE connections and xlates necessary to permit PPTP traffic. Only Version 1, as defined in RFC 2637, is supported.

PAT is only performed for the modified version of GRE **[**RFC 2637**]** when negotiated over the PPTP TCP control channel. Port Address Translation is *not* performed for the unmodified version of GRE **[**RFC 1701, RFC 1702**]**.

Specifically, the security appliance inspects the PPTP version announcements and the outgoing call request/response sequence. Only PPTP Version 1, as defined in RFC 2637, is inspected. Further inspection on the TCP control channel is disabled if the version announced by either side is not Version 1. In addition, the outgoing-call request and reply sequence are tracked. Connections and xlates are dynamic allocated as necessary to permit subsequent secondary GRE data traffic.

The PPTP inspection engine must be enabled for PPTP traffic to be translated by PAT. Additionally, PAT is only performed for a modified version of GRE (RFC2637) and only if it is negotiated over the PPTP TCP control channel. PAT is not performed for the unmodified version of GRE (RFC 1701 and RFC 1702).

As described in RFC 2637, the PPTP protocol is mainly used for the tunneling of PPP sessions initiated from a modem bank PAC (PPTP Access Concentrator) to the headend PNS (PPTP Network Server). When used this way, the PAC is the remote client and the PNS is the server.

However, when used for VPN by Windows, the interaction is inverted. The PNS is a remote single-user PC that initiates connection to the head-end PAC to gain access to a central network.

RADIUS Accounting Inspection

See the "Select RADIUS Accounting Map" section on page 23-16 for information about RADIUS accounting inspection.

RSH Inspection

RSH inspection is enabled by default. The RSH protocol uses a TCP connection from the RSH client to the RSH server on TCP port 514. The client and server negotiate the TCP port number where the client listens for the STDERR output stream. RSH inspection supports NAT of the negotiated port number if necessary.

RTSP Inspection

This section describes RTSP application inspection. This section includes the following topics:

- RTSP Inspection Overview, page 25-22
- Using RealPlayer, page 25-23
- Restrictions and Limitations, page 25-23

RTSP Inspection Overview

The RTSP inspection engine lets the security appliance pass RTSP packets. RTSP is used by RealAudio, RealNetworks, Apple QuickTime 4, RealPlayer, and Cisco IP/TV connections.

Note For Cisco IP/TV, use RTSP TCP port 554 and TCP 8554.

RTSP applications use the well-known port 554 with TCP (rarely UDP) as a control channel. The security appliance only supports TCP, in conformity with RFC 2326. This TCP control channel is used to negotiate the data channels that is used to transmit audio/video traffic, depending on the transport mode that is configured on the client.

The supported RDT transports are: rtp/avp, rtp/avp/udp, x-real-rdt, x-real-rdt/udp, and x-pn-tng/udp.

The security appliance parses Setup response messages with a status code of 200. If the response message is travelling inbound, the server is outside relative to the security appliance and dynamic channels need to be opened for connections coming inbound from the server. If the response message is outbound, then the security appliance does not need to open dynamic channels.

Because RFC 2326 does not require that the client and server ports must be in the SETUP response message, the security appliance keeps state and remembers the client ports in the SETUP message. QuickTime places the client ports in the SETUP message and then the server responds with only the server ports.

RTSP inspection does not support PAT or dual-NAT. Also, the security appliance cannot recognize HTTP cloaking where RTSP messages are hidden in the HTTP messages.

Using RealPlayer

When using RealPlayer, it is important to properly configure transport mode. For the security appliance, add an Access Rule from the server to the client or vice versa. For RealPlayer, change transport mode by choosing **Options**>**Preferences**>**Transport**>**RTSP Settings**.

If using TCP mode on the RealPlayer, check the **Use TCP to Connect to Server** and **Attempt to use TCP for all content** check boxes. On the security appliance, there is no need to configure the inspection engine.

If using UDP mode on the RealPlayer, check the **Use TCP to Connect to Server** and **Attempt to use UDP for static content** check boxes, and for live content not available via Multicast. On the security appliance, add an **inspect rtsp** *port* command.

Restrictions and Limitations

The following restrictions apply to RTSP inspection:

- The security appliance does not support multicast RTSP or RTSP messages over UDP.
- PAT is not supported.
- The security appliance does not have the ability to recognize HTTP cloaking where RTSP messages are hidden in the HTTP messages.
- The security appliance cannot perform NAT on RTSP messages because the embedded IP addresses are contained in the SDP files as part of HTTP or RTSP messages. Packets could be fragmented and security appliance cannot perform NAT on fragmented packets.
- With Cisco IP/TV, the number of translates the security appliance performs on the SDP part of the message is proportional to the number of program listings in the Content Manager (each program listing can have at least six embedded IP addresses).
- You can configure NAT for Apple QuickTime 4 or RealPlayer. Cisco IP/TV only works with NAT if the Viewer and Content Manager are on the outside network and the server is on the inside network.

SIP Inspection

This section describes SIP application inspection. This section includes the following topics:

- SIP Inspection Overview, page 25-24

- SIP Instant Messaging, page 25-24

SIP Inspection Overview

SIP, as defined by the IETF, enables call handling sessions, particularly two-party audio conferences, or “calls.” SIP works with SDP for call signalling. SDP specifies the ports for the media stream. Using SIP, the security appliance can support any SIP VoIP gateways and VoIP proxy servers. SIP and SDP are defined in the following RFCs:

- SIP: Session Initiation Protocol, RFC 2543
- SDP: Session Description Protocol, RFC 2327

To support SIP calls through the security appliance, signaling messages for the media connection addresses, media ports, and embryonic connections for the media must be inspected, because while the signaling is sent over a well-known destination port (UDP/TCP 5060), the media streams are dynamically allocated. Also, SIP embeds IP addresses in the user-data portion of the IP packet. SIP inspection applies NAT for these embedded IP addresses.

The following limitations and restrictions apply when using PAT with SIP:

- If a remote endpoint tries to register with a SIP proxy on a network protected by the security appliance, the registration fails under very specific conditions, as follows:
  - PAT is configured for the remote endpoint.
  - The SIP registrar server is on the outside network.
  - The port is missing in the contact field in the REGISTER message sent by the endpoint to the proxy server.
- If a SIP device transmits a packet in which the SDP portion has an IP address in the owner/creator field (o=) that is different than the IP address in the connection field (c=), the IP address in the o= field may not be properly translated. This is due to a limitation in the SIP protocol, which does not provide a port value in the o= field.

SIP Instant Messaging

Instant Messaging refers to the transfer of messages between users in near real-time. SIP supports the Chat feature on Windows XP using Windows Messenger RTC Client version 4.7.0105 only. The MESSAGE/INFO methods and 202 Accept response are used to support IM as defined in the following RFCs:

- Session Initiation Protocol (SIP)-Specific Event Notification, RFC 3265
- Session Initiation Protocol (SIP) Extension for Instant Messaging, RFC 3428

MESSAGE/INFO requests can come in at any time after registration/subscription. For example, two users can be online at any time, but not chat for hours. Therefore, the SIP inspection engine opens pinholes that time out according to the configured SIP timeout value. This value must be configured at least five minutes longer than the subscription duration. The subscription duration is defined in the Contact Expires value and is typically 30 minutes.

Because MESSAGE/INFO requests are typically sent using a dynamically allocated port other than port 5060, they are required to go through the SIP inspection engine.

Note Only the Chat feature is currently supported. Whiteboard, File Transfer, and Application Sharing are not supported. RTC Client 5.0 is not supported.

SIP inspection translates the SIP text-based messages, recalculates the content length for the SDP portion of the message, and recalculates the packet length and checksum. It dynamically opens media connections for ports specified in the SDP portion of the SIP message as address/ports on which the endpoint should listen.

SIP inspection has a database with indices CALL_ID/FROM/TO from the SIP payload. These indices identify the call, the source, and the destination. This database contains the media addresses and media ports found in the SDP media information fields and the media type. There can be multiple media addresses and ports for a session. The security appliance opens RTP/RTCP connections between the two endpoints using these media addresses/ports.

The well-known port 5060 must be used on the initial call setup (INVITE) message; however, subsequent messages may not have this port number. The SIP inspection engine opens signaling connection pinholes, and marks these connections as SIP connections. This is done for the messages to reach the SIP application and be translated.

As a call is set up, the SIP session is in the "transient" state until the media address and media port is received from the called endpoint in a Response message indicating the RTP port the called endpoint listens on. If there is a failure to receive the response messages within one minute, the signaling connection is torn down.

Once the final handshake is made, the call state is moved to active and the signaling connection remains until a BYE message is received.

If an inside endpoint initiates a call to an outside endpoint, a media hole is opened to the outside interface to allow RTP/RTCP UDP packets to flow to the inside endpoint media address and media port specified in the INVITE message from the inside endpoint. Unsolicited RTP/RTCP UDP packets to an inside interface does not traverse the security appliance, unless the security appliance configuration specifically allows it.

Skinny (SCCP) Inspection

This section describes SCCP application inspection. This section includes the following topics:

- SCCP Inspection Overview, page 25-26
- Supporting Cisco IP Phones, page 25-26
- Restrictions and Limitations, page 25-27

Note For specific information about setting up the Phone Proxy on the security appliance, which is part of the Cisco Unified Communications architecture and supports IP phone deployment, see the "Phone Proxy" section on page 25-142.

SCCP Inspection Overview

Skinny (SCCP) is a simplified protocol used in VoIP networks. Cisco IP Phones using SCCP can coexist in an H.323 environment. When used with Cisco CallManager, the SCCP client can interoperate with H.323 compliant terminals. Application layer functions in the security appliance recognize SCCP Version 3.3. There are 5 versions of the SCCP protocol: 2.4, 3.0.4, 3.1.1, 3.2, and 3.3.2. The security appliance supports all versions through Version 3.3.2.

Note For specific information about setting up the Phone Proxy on the security appliance, which is part of the Cisco Unified Communications architecture and supports IP phone deployment, see the "Phone Proxy" section on page 25-142.

The security appliance supports PAT and NAT for SCCP. PAT is necessary if you have more IP phones than global IP addresses for the IP phones to use. By supporting NAT and PAT of SCCP Signaling packets, Skinny application inspection ensures that all SCCP signalling and media packets can traverse the security appliance.

Normal traffic between Cisco CallManager and Cisco IP Phones uses SCCP and is handled by SCCP inspection without any special configuration. The security appliance also supports DHCP options 150 and 66, which it accomplishes by sending the location of a TFTP server to Cisco IP Phones and other DHCP clients. Cisco IP Phones might also include DHCP option 3 in their requests, which sets the default route.

Supporting Cisco IP Phones

Note For specific information about setting up the Phone Proxy on the security appliance, which is part of the Cisco Unified Communications architecture and supports IP phone deployment, see see the "Phone Proxy" section on page 25-142.

In topologies where Cisco CallManager is located on the higher security interface with respect to the Cisco IP Phones, if NAT is required for the Cisco CallManager IP address, the mapping must be **static** as a Cisco IP Phone requires the Cisco CallManager IP address to be specified explicitly in its configuration. An static identity entry allows the Cisco CallManager on the higher security interface to accept registrations from the Cisco IP Phones.

Cisco IP Phones require access to a TFTP server to download the configuration information they need to connect to the Cisco CallManager server.

When the Cisco IP Phones are on a lower security interface compared to the TFTP server, you must use an access list to connect to the protected TFTP server on UDP port 69. While you do need a static entry for the TFTP server, this does not have to be an identity static entry. When using NAT, an identity static entry maps to the same IP address. When using PAT, it maps to the same IP address and port.

When the Cisco IP Phones are on a *higher* security interface compared to the TFTP server and Cisco CallManager, no access list or static entry is required to allow the Cisco IP Phones to initiate the connection.

Restrictions and Limitations

Note For specific information about setting up the Phone Proxy on the security appliance, which is part of the Cisco Unified Communications architecture and supports IP phone deployment, see the "Phone Proxy" section on page 25-142.

The following are limitations that apply to the current version of PAT and NAT support for SCCP:

- PAT does not work with configurations containing the **alias** command.
- Outside NAT or PAT is *not* supported.

If the address of an internal Cisco CallManager is configured for NAT or PAT to a different IP address or port, registrations for external Cisco IP Phones fail because the security appliance currently does not support NAT or PAT for the file content transferred over TFTP. Although the security appliance supports NAT of TFTP messages and opens a pinhole for the TFTP file, the security appliance cannot translate the Cisco CallManager IP address and port embedded in the Cisco IP Phone configuration files that are transferred by TFTP during phone registration.

Note The security appliance supports Stateful Failover of SCCP calls except for calls that are in the middle of call setup.

SMTP and Extended SMTP Inspection

ESMTP application inspection provides improved protection against SMTP-based attacks by restricting the types of SMTP commands that can pass through the security appliance and by adding monitoring capabilities.

ESMTP is an enhancement to the SMTP protocol and is similar is most respects to SMTP. For convenience, the term SMTP is used in this document to refer to both SMTP and ESMTP. The application inspection process for extended SMTP is similar to SMTP application inspection and includes support for SMTP sessions. Most commands used in an extended SMTP session are the same as those used in an SMTP session but an ESMTP session is considerably faster and offers more options related to reliability and security, such as delivery status notification.

Extended SMTP application inspection adds support for eight extended SMTP commands, including AUTH, EHLO, ETRN, HELP, SAML, SEND, SOML and VRFY. Along with the support for seven RFC 821 commands (DATA, HELO, MAIL, NOOP, QUIT, RCPT, RSET), the security appliance supports a total of fifteen SMTP commands.

Other extended SMTP commands, such as ATRN, STARTLS, ONEX, VERB, CHUNKING, and private extensions and are not supported. Unsupported commands are translated into Xs, which are rejected by the internal server. This results in a message such as "500 Command unknown: 'XXX'." Incomplete commands are discarded.

The ESMTP inspection engine changes the characters in the server SMTP banner to asterisks except for the "2", "0", "0" characters. Carriage return (CR) and linefeed (LF) characters are ignored.

With SMTP inspection enabled, a Telnet session used for interactive SMTP may hang if the following rules are not observed: SMTP commands must be at least four characters in length; must be terminated with carriage return and line feed; and must wait for a response before issuing the next reply.

An SMTP server responds to client requests with numeric reply codes and optional human-readable strings. SMTP application inspection controls and reduces the commands that the user can use as well as the messages that the server returns. SMTP inspection performs three primary tasks:

- Restricts SMTP requests to seven basic SMTP commands and eight extended commands.
- Monitors the SMTP command-response sequence.
- Generates an audit trail—Audit record 108002 is generated when invalid character embedded in the mail address is replaced. For more information, see RFC 821.

SMTP inspection monitors the command and response sequence for the following anomalous signatures:

- Truncated commands.
- Incorrect command termination (not terminated with <CR><LR>).
- The MAIL and RCPT commands specify who are the sender and the receiver of the mail. Mail addresses are scanned for strange characters. The pipeline character (|) is deleted (changed to a blank space) and "<" ,">" are only allowed if they are used to define a mail address (">" must be preceded by "<").
- Unexpected transition by the SMTP server.
- For unknown commands, the security appliance changes all the characters in the packet to X. In this case, the server generates an error code to the client. Because of the change in the packed, the TCP checksum has to be recalculated or adjusted.
- TCP stream editing.
- Command pipelining.

SNMP Inspection

SNMP application inspection lets you restrict SNMP traffic to a specific version of SNMP. Earlier versions of SNMP are less secure; therefore, denying certain SNMP versions may be required by your security policy. The security appliance can deny SNMP Versions 1, 2, 2c, or 3. You control the versions permitted by creating an SNMP map.

SQL*Net Inspection

SQL*Net inspection is enabled by default.

The SQL*Net protocol consists of different packet types that the security appliance handles to make the data stream appear consistent to the Oracle applications on either side of the security appliance.

The default port assignment for SQL*Net is 1521. This is the value used by Oracle for SQL*Net, but this value does not agree with IANA port assignments for Structured Query Language (SQL). Use the **class-map** command to apply SQL*Net inspection to a range of port numbers.

The security appliance translates all addresses and looks in the packets for all embedded ports to open for SQL*Net Version 1.

For SQL*Net Version 2, all DATA or REDIRECT packets that immediately follow REDIRECT packets with a zero data length will be fixed up.

The packets that need fix-up contain embedded host/port addresses in the following format:

```
(ADDRESS=(PROTOCOL=tcp)(DEV=6)(HOST=a.b.c.d)(PORT=a))
```

SQL*Net Version 2 TNSFrame types (Connect, Accept, Refuse, Resend, and Marker) will not be scanned for addresses to NAT nor will inspection open dynamic connections for any embedded ports in the packet.

SQL*Net Version 2 TNSFrames, Redirect, and Data packets will be scanned for ports to open and addresses to NAT, if preceded by a REDIRECT TNSFrame type with a zero data length for the payload. When the Redirect message with data length zero passes through the security appliance, a flag will be set in the connection data structure to expect the Data or Redirect message that follows to be translated and ports to be dynamically opened. If one of the TNS frames in the preceding paragraph arrive after the Redirect message, the flag will be reset.

The SQL*Net inspection engine will recalculate the checksum, change IP, TCP lengths, and readjust Sequence Numbers and Acknowledgment Numbers using the delta of the length of the new and old message.

SQL*Net Version 1 is assumed for all other cases. TNSFrame types (Connect, Accept, Refuse, Resend, Marker, Redirect, and Data) and all packets will be scanned for ports and addresses. Addresses will be translated and port connections will be opened.

Sun RPC Inspection

This section describes Sun RPC application inspection. This section includes the following topics:

- Sun RPC Inspection Overview, page 25-29
- SUNRPC Server, page 25-29

Sun RPC Inspection Overview

The Sun RPC inspection engine enables or disables application inspection for the Sun RPC protocol. Sun RPC is used by NFS and NIS. Sun RPC services can run on any port. When a client attempts to access an Sun RPC service on a server, it must learn the port that service is running on. It does this by querying the port mapper process, usually rpcbind, on the well-known port of 111.

The client sends the Sun RPC program number of the service and the port mapper process responds with the port number of the service. The client sends its Sun RPC queries to the server, specifying the port identified by the port mapper process. When the server replies, the security appliance intercepts this packet and opens both embryonic TCP and UDP connections on that port.

Note NAT or PAT of Sun RPC payload information is not supported.

SUNRPC Server

The Configuration > Firewall > Advanced > SUNRPC Server pane 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

The Configuration > Firewall > Advanced > SUNRPC Server > 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
•	•	•	•	—

TFTP Inspection

TFTP inspection is enabled by default.

TFTP, described in RFC 1350, is a simple protocol to read and write files between a TFTP server and client.

The security appliance inspects TFTP traffic and dynamically creates connections and translations, if necessary, to permit file transfer between a TFTP client and server. Specifically, the inspection engine inspects TFTP read request (RRQ), write request (WRQ), and error notification (ERROR).

A dynamic secondary channel and a PAT translation, if necessary, are allocated on a reception of a valid read (RRQ) or write (WRQ) request. This secondary channel is subsequently used by TFTP for file transfer or error notification.

Only the TFTP server can initiate traffic over the secondary channel, and at most one incomplete secondary channel can exist between the TFTP client and server. An error notification from the server closes the secondary channel.

TFTP inspection must be enabled if static PAT is used to redirect TFTP traffic.

XDMCP Inspection

XDMCP inspection is enabled by default; however, the XDMCP inspection engine is dependent upon proper configuration of the **established** command.

XDMCP is a protocol that uses UDP port 177 to negotiate X sessions, which use TCP when established.

For successful negotiation and start of an XWindows session, the security appliance must allow the TCP back connection from the Xhosted computer. To permit the back connection, use the **established** command on the security appliance. Once XDMCP negotiates the port to send the display, The **established** command is consulted to verify if this back connection should be permitted.

During the XWindows session, the manager talks to the display Xserver on the well-known port 6000 | n. Each display has a separate connection to the Xserver, as a result of the following terminal setting.

```
setenv DISPLAY Xserver:n
```

where *n* is the display number.

When XDMCP is used, the display is negotiated using IP addresses, which the security appliance can NAT if needed. XDCMP inspection does not support PAT.

Service Policy Field Descriptions

This section lists the field decsriptions for each protocol inspection dialog box, and includes the following topics:

- Rule Actions > Protocol Inspection Tab, page 25-32
- Select DCERPC Map, page 25-34
- Select DNS Map, page 25-34
- Select ESMTP Map, page 25-35
- Select FTP Map, page 25-35
- Select GTP Map, page 25-36
- Select H.323 Map, page 25-36
- Select HTTP Map, page 25-37

- Select IM Map, page 25-37
- Select IP Options Inspect Map, page 25-38
- Select IPSec-Pass-Thru Map, page 25-38
- Select MGCP Map, page 25-39
- Select NETBIOS Map, page 25-39
- Select RTSP Map, page 25-40
- Select SCCP (Skinny) Map, page 25-40
- Select SIP Map, page 25-41
- Select SNMP Map, page 25-42

Rule Actions > Protocol Inspection Tab

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.
- IP-Options**—Enables inspection for the presence of IP options in an IP packet.**
  - Configure—Displays the Select IP Options Map dialog box, which lets you select the default map to use or add a map for this inspection.
- 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

Inspect Map Field Descriptions, page 25-64

Inspect command pages for each protocol in the *Cisco ASA 5500 Series Command Reference*.

Select DCERPC Map

The Select DCERPC Map dialog box lets you select or create a new DCERPC map. A DCERPC map lets you change the configuration values used for DCERPC application inspection. The Select DCERPC Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default DCERPC inspection map—Specifies to use the default DCERPC map.
- Select a DCERPC map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

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

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

- Use the default DNS inspection map—Specifies to use the default 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

The Select ESMTP Map dialog box lets you select or create a new ESMTP map. An ESMTP map lets you change the configuration values used for ESMTP application inspection. The Select ESMTP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default ESMTP inspection map—Specifies to use the default ESMTP map.
- Select an ESMTP map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

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

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.
- Use the default FTP inspection map—Specifies to use the default FTP map.
- Select an FTP map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

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

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

- Use the default GTP inspection map—Specifies to use the default GTP map.
- Select an GTP map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

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

The Select H.323 Map dialog box lets you select or create a new H.323 map. An H.323 map lets you change the configuration values used for H.323 application inspection. The Select H.323 Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default H.323 inspection map—Specifies to use the default H.323 map.
- Select an H.323 map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

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

The Select HTTP Map dialog box lets you select or create a new HTTP map. An HTTP map lets you change the configuration values used for HTTP application inspection. The Select HTTP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default HTTP inspection map—Specifies to use the default HTTP map.
- Select an HTTP map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

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

The Select IM Map dialog box lets you select or create a new IM map. An IM map lets you change the configuration values used for IM application inspection. The Select IM Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Add—Opens the Add Policy Map dialog box for the inspection.

Modes

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

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

Select IP Options Inspect Map

The Select IP-Options Inspect Map dialog box lets you select or create a new IP Options inspection map. Use this inspection map to control whether the security appliance drops, passes, or clears IP packets containing the following IP options—End of Options List, No Operations, and Router Alert.

Fields

- Use the default IP-Options inspection map—Specifies to use the default IP Options map. The default map drops packets containing all the inspected IP options, namely End of Options List (EOOL), No Operation (NOP), and Router Alert (RTRALT).
- Select an IP-Options inspect map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add IP-Options Inspect Map dialog box for the inspection.

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

The Select IPSec-Pass-Thru dialog box lets you select or create a new IPSec map. An IPSec map lets you change the configuration values used for IPSec application inspection. The Select IPSec Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default IPSec inspection map—Specifies to use the default IPSec map.
- Select an IPSec map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

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

The Select MGCP Map dialog box lets you select or create a new MGCP map. An MGCP map lets you change the configuration values used for MGCP application inspection. The Select MGCP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default MGCP inspection map—Specifies to use the default MGCP map.
- Select an MGCP map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

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

The Select NETBIOS Map dialog box lets you select or create a new NetBIOS map. A NetBIOS map lets you change the configuration values used for NetBIOS application inspection. The Select NetBIOS Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default IM inspection map—Specifies to use the default NetBIOS map.
- Select a NetBIOS map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

Modes

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

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

Select RTSP Map

The Select RTSP Map dialog box lets you select or create a new RTSP map. An RTSP map lets you change the configuration values used for RTSP application inspection. The Select RTSP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default RTSP inspection map—Specifies to use the default RTSP inspection map.
- Select a RTSP inspect map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

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

The Select SCCP (Skinny) Map dialog box lets you select or create a new SCCP (Skinny) map. An SCCP (Skinny) map lets you change the configuration values used for SCCP (Skinny) application inspection. The Select SCCP (Skinny) Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default SCCP (Skinny) inspection map—Specifies to use the default SCCP (Skinny) map.
- Select an SCCP (Skinny) map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.
- Encrypted Traffic Inspection—Lets you specify TLS proxy settings for the inspect map.
  - Do not inspect Encrypted Traffic—Disables the inspection of Skinny application inspection.
  - Use Phone Proxy to enable inspection of encrypted traffic—Uses the Phone Proxy configured on the security appliance to inspect Skinny application traffic. See the “Phone Proxy” section on page 25-142.

  - Use TLS Proxy to enable inspection of encrypted traffic—Specifies to use Transaction Layer Security Proxy to enable inspection of encryped traffic.

    TLS Proxy Name:—Name of existing TLS Proxy.

    New—Opens the Add TLS Proxy dialog box to add a TLS Proxy.

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

The Select SIP Map dialog box lets you select or create a new SIP map. A SIP map lets you change the configuration values used for SIP application inspection. The Select SIP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default SIP inspection map—Specifies to use the default SIP map.
- Select a SIP map for fine control over inspection**—**Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.
- TLS Proxy—Lets you specify TLS proxy settings for the inspect map.
  - Use TLS Proxy to enable inspection of encrypted traffic—Specifies to use Transaction Layer Security Proxy to enable inspection of encryped traffic.

    TLS Proxy Name:—Name of existing TLS Proxy.

    New—Opens the Add TLS Proxy dialog box to add a TLS Proxy.

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

The Select SNMP Map dialog box lets you select or create a new SNMP map. An SNMP map lets you change the configuration values used for SNMP application inspection. The Select SNMP Map table provides a list of previously configured maps that you can select for application inspection.

Fields

- Use the default SNMP inspection map—Specifies to use the default SNMP map.
- Select an SNMP map for fine control over inspection—Lets you select a defined application inspection map or add a new one.
- Add—Opens the Add Policy Map dialog box for the inspection.

Modes

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

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

Class Map Field Descriptions

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, page 25-42
- FTP Class Map, page 25-47
- H.323 Class Map, page 25-50
- HTTP Class Map, page 25-52
- IM Class Map, page 25-57
- SIP Class Map, page 25-60

DNS Class Map

The DNS Class Map pane 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

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

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

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

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
•	•	•	•	—

Manage 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
•	•	•	•	—

FTP Class Map

The FTP Class Map pane lets you configure FTP class maps for FTP inspection.

An inspection class map matches application traffic with criteria specific to the application. You then identify the class map in the inspect map and enable actions. The difference between creating a class map and defining the traffic match directly in the inspect map is that you can create more complex match criteria and you can reuse class maps. The applications that support inspection class maps are DNS, FTP, H.323, HTTP, IM, and SIP.

Fields

- Name—Shows the FTP class map name.
- Match Conditions—Shows the type, match criterion, and value in the class map.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the FTP class map.
  - Value—Shows the value to match in the FTP class map.
- Description—Shows the description of the class map.
- Add—Adds an FTP class map.

- Edit—Edits an FTP class map.
- Delete—Deletes an FTP class map.

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

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

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

The H.323 Class Map pane lets you configure H.323 class maps for H.323 inspection.

An inspection class map matches application traffic with criteria specific to the application. You then identify the class map in the inspect map and enable actions. The difference between creating a class map and defining the traffic match directly in the inspect map is that you can create more complex match criteria and you can reuse class maps. The applications that support inspection class maps are DNS, FTP, H.323, HTTP, IM, and SIP.

Fields

- Name—Shows the H.323 class map name.
- Match Conditions—Shows the type, match criterion, and value in the class map.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the H.323 class map.
  - Value—Shows the value to match in the H.323 class map.
- Description—Shows the description of the class map.
- Add—Adds an H.323 class map.
- Edit—Edits an H.323 class map.
- Delete—Deletes an H.323 class map.

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

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

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

The HTTP Class Map pane lets you configure HTTP class maps for HTTP inspection.

An inspection class map matches application traffic with criteria specific to the application. You then identify the class map in the inspect map and enable actions. The difference between creating a class map and defining the traffic match directly in the inspect map is that you can create more complex match criteria and you can reuse class maps. The applications that support inspection class maps are DNS, FTP, H.323, HTTP, IM, and SIP.

Fields

- Name—Shows the HTTP class map name.
- Match Conditions—Shows the type, match criterion, and value in the class map.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the HTTP class map.
  - Value—Shows the value to match in the HTTP class map.
- Description—Shows the description of the class map.
- Add—Adds an HTTP class map.

- Edit—Edits an HTTP class map.
- Delete—Deletes an HTTP class map.

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

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

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

The IM Class Map pane 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

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

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

The SIP Class Map pane lets you configure SIP class maps for SIP inspection.

An inspection class map matches application traffic with criteria specific to the application. You then identify the class map in the inspect map and enable actions. The difference between creating a class map and defining the traffic match directly in the inspect map is that you can create more complex match criteria and you can reuse class maps. The applications that support inspection class maps are DNS, FTP, H.323, HTTP, IM, and SIP.

Fields

- Name—Shows the SIP class map name.
- Match Conditions—Shows the type, match criterion, and value in the class map.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the SIP class map.
  - Value—Shows the value to match in the SIP class map.
- Description—Shows the description of the class map.
- Add—Adds a SIP class map.
- Edit—Edits a SIP class map.
- Delete—Deletes a SIP class map.

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

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

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
•	•	•	•	—

Inspect Map Field Descriptions

This section describes how to configure inspect maps, and includes the following topics:

Note For information about RADIUS inspect maps, see the "Adding a Service Policy Rule for Management Traffic" section on page 23-11.

- DCERPC Inspect Map, page 25-67
- DNS Inspect Map, page 25-69
- ESMTP Inspect Map, page 25-75
- FTP Inspect Map, page 25-83
- GTP Inspect Map, page 25-87
- H.323 Inspect Map, page 25-93
- HTTP Inspect Map, page 25-99
- Instant Messaging (IM) Inspect Map, page 25-107
- IPSec Pass Through Inspect Map, page 25-110
- IP Options Inspect Map, page 25-113
- MGCP Inspect Map, page 25-114
- NetBIOS Inspect Map, page 25-117
- RTSP Inspect Map, page 25-118
- SCCP (Skinny) Inspect Map, page 25-120
- SIP Inspect Map, page 25-125
- SNMP Inspect Map, page 25-131

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.
IP Options	You can configure IP Options inspection to control which IP packets with specific IP options are allowed through the security appliance. Configuring this inspection instructs the security appliance to allow a packet to pass or to clear the specified IP options and then allow the packet to pass.
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.
RTSP	The RTSP inspection lets you create, view, and manage RTSP inspect maps. You can use an RTSP map to protect RTSP traffic, including RTSP PAT.
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

The DCERPC pane lets you view previously configured DCERPC application inspection maps. A DCERPC map lets you change the default configuration values used for DCERPC application inspection.

DCERPC is a protocol widely used by Microsoft distributed client and server applications that allows software clients to execute programs on a server remotely.

This typically involves a client querying a server called the Endpoint Mapper (EPM) listening on a well known port number for the dynamically allocated network information of a required service. The client then sets up a secondary connection to the server instance providing the service. The security appliance allows the appropriate port number and network address and also applies NAT, if needed, for the secondary connection.

DCERPC inspect maps inspect for native TCP communication between the EPM and client on well known TCP port 135. Map and lookup operations of the EPM are supported for clients. Client and server can be located in any security zone. The embedded server IP address and Port number are received from the applicable EPM response messages. Because a client may attempt multiple connections to the server port returned by EPM, multiple use of pinholes are allowed, which have user configurable timeouts.

Fields

- DCERPC Inspect Maps—Table that lists the defined DCERPC inspect maps.
- Add—Configures a new DCERPC inspect map. To edit a DCERPC inspect map, choose the DCERPC entry in the DCERPC Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the DCERPC Inspect Maps table.
- Security Level—Select the security level (high, medium, or low).
  - Low

    Pinhole timeout: 00:02:00

    Endpoint mapper service: not enforced

    Endpoint mapper service lookup: enabled

    Endpoint mapper service lookup timeout: 00:05:00
  - Medium—Default.

    Pinhole timeout: 00:01:00

    Endpoint mapper service: not enforced

    Endpoint mapper service lookup: disabled.
  - High

    Pinhole timeout: 00:01:00

    Endpoint mapper service: enforced

    Endpoint mapper service lookup: disabled
  - Customize—Opens the Add/Edit DCERPC Policy Map dialog box for additional settings.
  - Default Level—Sets the security level back to the default level of Medium.

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 DCERPC Policy Map

The Add/Edit DCERPC Policy Map pane lets you configure the security level and parameters for DCERPC application inspection maps.

Fields

- Name—When adding a DCERPC map, enter the name of the DCERPC map. When editing a DCERPC map, the name of the previously configured DCERPC map is shown.
- Description—Enter the description of the DCERPC map, up to 200 characters in length.
- Security Level—Select the security level (high, medium, or low).
  - Low

    Pinhole timeout: 00:02:00

    Endpoint mapper service: not enforced

    Endpoint mapper service lookup: enabled

    Endpoint mapper service lookup timeout: 00:05:00
  - Medium—Default.

    Pinhole timeout: 00:01:00

    Endpoint mapper service: not enforced

    Endpoint mapper service lookup: disabled.
  - High

    Pinhole timeout: 00:01:00

    Endpoint mapper service: enforced

    Endpoint mapper service lookup: disabled
  - Default Level—Sets the security level back to the default level of Medium.
- Details—Shows the Parameters to configure additional settings.
  - Pinhole Timeout—Sets the pinhole timeout. Because a client may use the server information returned by the endpoint mapper for multiple connections, the timeout value is configurable based on the client application environment. Range is from 0:0:1 to 1193:0:0. Default is 2 minutes.
  - Enforce endpoint-mapper service—Enforces endpoint mapper service during binding.
  - Enable endpoint-mapper service lookup—Enables the lookup operation of the endpoint mapper service. If disabled, the pinhole timeout is used.

    Enforce Service Lookup Timeout—Enforces the service lookup timeout specified.

    Service Lookup Timeout—Sets the timeout for pinholes from lookup operation.

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

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

- DNS Inspect Maps—Table that lists the defined DNS inspect maps.
- Add—Configures a new DNS inspect map. To edit a DNS inspect map, choose the DNS entry in the DNS Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the DNS Inspect Maps table.
- 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 Add/Edit DNS Policy Map 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
•	•	•	•	—

Add/Edit DNS Policy Map (Security Level)

The Add/Edit DNS Policy Map pane lets you configure the security level and additional settings for DNS application inspection maps.

Fields

- Name—When adding a DNS map, enter the name of the DNS map. When editing a DNS map, the name of the previously configured DNS map is shown.
- Description—Enter the description of the DNS 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

- Default Level—Sets the security level back to the default level of Low.

- Details—Shows the Protocol Conformance, Filtering, Mismatch Rate, and Inspection tabs to configure additional settings.

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 Policy Map (Details)

The Add/Edit DNS Policy Map pane lets you configure the security level and additional settings for DNS application inspection maps

Fields

- Name—When adding a DNS map, enter the name of the DNS map. When editing a DNS map, the name of the previously configured DNS map is shown.
- Description—Enter the description of the DNS map, up to 200 characters in length.
- Security Level—Shows the security level to configure.
- 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

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
•	•	•	•	—

ESMTP Inspect Map

The ESMTP pane lets you view previously configured ESMTP application inspection maps. An ESMTP map lets you change the default configuration values used for ESMTP application inspection.

Since ESMTP traffic can be a main source of attack from spam, phising, malformed messages, buffer overflows, and buffer underflows, detailed packet inspection and control of ESMTP traffic are supported. Application security and protocol conformance enforce the sanity of the ESMTP message as well as detect several attacks, block senders and receivers, and block mail relay.

Fields

- ESMTP Inspect Maps—Table that lists the defined ESMTP inspect maps.
- Add—Configures a new ESMTP inspect map. To edit an ESMTP inspect map, choose the ESMTP entry in the ESMTP Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the ESMTP Inspect Maps table.
- Security Level—Select the security level (high, medium, or low).
  - Low—Default.

    Log if command line length is greater than 512

    Log if command recipient count is greater than 100

    Log if body line length is greater than 1000

    Log if sender address length is greater than 320

    Log if MIME file name length is greater than 255
  - Medium

Obfuscate Server Banner

Drop Connections if command line length is greater than 512

Drop Connections if command recipient count is greater than 100

Drop Connections if body line length is greater than 1000

Drop Connections if sender address length is greater than 320

Drop Connections if MIME file name length is greater than 255

- High

  Obfuscate Server Banner

  Drop Connections if command line length is greater than 512

  Drop Connections if command recipient count is greater than 100

  Drop Connections if body line length is greater than 1000

  Drop Connections and log if sender address length is greater than 320

  Drop Connections and log if MIME file name length is greater than 255
- MIME File Type Filtering—Opens the MIME Type Filtering dialog box to configure MIME file type filters.
- Customize—Opens the Add/Edit ESMTP Policy Map dialog box for additional settings.
- Default Level—Sets the security level back to the default level of Low.

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

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
•	•	•	•	—

Add/Edit ESMTP Policy Map (Security Level)

The Add/Edit ESMTP Policy Map pane lets you configure the security level and additional settings for ESMTP application inspection maps.

Fields

- Name—When adding an ESMTP map, enter the name of the ESMTP map. When editing an ESMTP map, the name of the previously configured ESMTPS map is shown.
- Description—Enter the description of the ESMTP map, up to 200 characters in length.
- Security Level—Select the security level (high, medium, or low).
  - Low—Default.

    Log if command line length is greater than 512

    Log if command recipient count is greater than 100

    Log if body line length is greater than 1000

    Log if sender address length is greater than 320

    Log if MIME file name length is greater than 255
  - Medium

    Obfuscate Server Banner

    Drop Connections if command line length is greater than 512

    Drop Connections if command recipient count is greater than 100

    Drop Connections if body line length is greater than 1000

    Drop Connections if sender address length is greater than 320

    Drop Connections if MIME file name length is greater than 255
  - High

    Obfuscate Server Banner

    Drop Connections if command line length is greater than 512

    Drop Connections if command recipient count is greater than 100

    Drop Connections if body line length is greater than 1000

    Drop Connections and log if sender address length is greater than 320

    Drop Connections and log if MIME file name length is greater than 255

  - MIME File Type Filtering—Opens the MIME Type Filtering dialog box to configure MIME file type filters.
  - Default Level—Sets the security level back to the default level of Low.
- Details—Shows the Parameters and Inspections tabs to configure additional settings.

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 Policy Map (Details)

The Add/Edit ESMTP Policy Map pane lets you configure the security level and additional settings for ESMTP application inspection maps.

Fields

- Name—When adding an ESMTP map, enter the name of the ESMTP map. When editing an ESMTP map, the name of the previously configured ESMTP map is shown.
- Description—Enter the description of the ESMTP map, up to 200 characters in length.
- Security Level—Shows the security level and mime file type filtering settings to configure.
- Parameters—Tab that lets you configure the parameters for the ESMTP inspect map.
  - Mask server banner—Enforces banner obfuscation.
  - Configure Mail Relay—Enables ESMTP mail relay.

    Domain Name—Specifies a local domain.

    Action—Drop connection or log.

    Log—Enable or disable.
- Inspections—Tab that shows you the ESMTP inspection configuration and lets you add or edit.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the ESMTP inspection.
  - Value—Shows the value to match in the ESMTP inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add ESMTP Inspect dialog box to add an ESMTP inspection.
  - Edit—Opens the Edit ESMTP Inspect dialog box to edit an ESMTP inspection.
  - Delete—Deletes an ESMTP inspection.
  - Move Up—Moves an inspection up in the list.
  - Move Down—Moves an inspection down in the list.

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

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

The FTP pane lets you view previously configured FTP application inspection maps. An FTP map lets you change the default configuration values used for FTP application inspection.

FTP command filtering and security checks are provided using strict FTP inspection for improved security and control. Protocol conformance includes packet length checks, delimiters and packet format checks, command terminator checks, and command validation.

Blocking FTP based on user values is also supported so that it is possible for FTP sites to post files for download, but restrict access to certain users. You can block FTP connections based on file type, server name, and other attributes. System message logs are generated if an FTP connection is denied after inspection.

Fields

- FTP Inspect Maps—Table that lists the defined FTP inspect maps.
- Add—Configures a new FTP inspect map. To edit an FTP inspect map, choose the FTP entry in the FTP Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the FTP Inspect Maps table.
- Security Level—Select the security level (medium or low).
  - Low

    Mask Banner Disabled

    Mask Reply Disabled
  - Medium—Default.

    Mask Banner Enabled

    Mask Reply Enabled
  - File Type Filtering—Opens the Type Filtering dialog box to configure file type filters.
  - Customize—Opens the Add/Edit FTP Policy Map dialog box for additional settings.
  - Default Level—Sets the security level back to the default level of Medium.

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

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
•	•	•	•	—

Add/Edit FTP Policy Map (Security Level)

The Add/Edit FTP Policy Map pane lets you configure the security level and additional settings for FTP application inspection maps.

Fields

- Name—When adding an FTP map, enter the name of the FTP map. When editing an FTP map, the name of the previously configured FTP map is shown.
- Description—Enter the description of the FTP map, up to 200 characters in length.
- Security Level—Select the security level (medium or low).
  - Low

    Mask Banner Disabled

    Mask Reply Disabled
  - Medium—Default.

    Mask Banner Enabled

    Mask Reply Enabled
  - File Type Filtering—Opens the Type Filtering dialog box to configure file type filters.
  - Default Level—Sets the security level back to the default level of Medium.
- Details—Shows the Parameters and Inspections tabs to configure additional settings.

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 Policy Map (Details)

The Add/Edit FTP Policy Map pane lets you configure the security level and additional settings for FTP application inspection maps.

Fields

- Name—When adding an FTP map, enter the name of the FTP map. When editing an FTP map, the name of the previously configured FTP map is shown.
- Description—Enter the description of the FTP map, up to 200 characters in length.
- Security Level—Shows the security level and file type filtering settings to configure.
- Parameters—Tab that lets you configure the parameters for the FTP inspect map.
  - Mask greeting banner from the server—Masks the greeting banner from the FTP server to prevent the client from discovering server information.
  - Mask reply to SYST command—Masks the reply to the syst command to prevent the client from discovering server information.
- Inspections—Tab that shows you the FTP inspection configuration and lets you add or edit.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the FTP inspection.
  - Value—Shows the value to match in the FTP inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add FTP Inspect dialog box to add an FTP inspection.
  - Edit—Opens the Edit FTP Inspect dialog box to edit an FTP inspection.
  - Delete—Deletes an FTP inspection.
  - Move Up—Moves an inspection up in the list.
  - Move Down—Moves an inspection down in the list.

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

The Add/Edit FTP Inspect dialog box lets you define the match criterion and value for the FTP inspect map.

Fields

- Single Match—Specifies that the FTP inspect has only one match statement.
- Match Type—Specifies whether traffic should match or not match the values.

  For example, if No Match is selected on the string “example.com,” then any traffic that contains “example.com” is excluded from the class map.
- Criterion—Specifies which criterion of FTP traffic to match.
  - Request Command—Match an FTP request command.
  - File Name—Match a filename for FTP transfer.
  - File Type—Match a file type for FTP transfer.
  - Server—Match an FTP server.
  - User Name—Match an FTP user.
- Request Command Criterion Values—Specifies the value details for FTP request command match.
  - Request Command:

    APPE—Command that appends to a file.

    CDUP—Command that changes to the parent directory of the current working directory.

    DELE—Command that deletes a file.

    GET—Command that gets a file.

    HELP—Command that provides help information.

    MKD—Command that creates a directory.

    PUT—Command that sends a file.

    RMD—Command that deletes a directory.

    RNFR—Command that specifies rename-from filename.

    RNTO—Command that specifies rename-to filename.

    SITE—Commands that are specific to the server system. Usually used for remote administration.

    STOU—Command that stores a file using a unique filename.
- File Name Criterion Values—Specifies the value details for FTP filename match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- File Type Criterion Values—Specifies the value details for FTP file type match.
  - Regular Expression—Lists the defined regular expressions to match.

  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Server Criterion Values—Specifies the value details for FTP server match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- User Name Criterion Values—Specifies the value details for FTP user name match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- Multiple Matches—Specifies multiple matches for the FTP inspection.
  - FTP Traffic Class—Specifies the FTP traffic class match.
  - Manage—Opens the Manage FTP Class Maps dialog box to add, edit, or delete FTP Class Maps.
- Action—Reset.
- Log—Enable or disable.

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

The GTP pane lets you view previously configured GTP application inspection maps. A GTP map lets you change the default configuration values used for GTP application inspection.

GTP is a relatively new protocol designed to provide security for wireless connections to TCP/IP networks, such as the Internet. You can use a GTP map to control timeout values, message sizes, tunnel counts, and GTP versions traversing the security appliance.

Note GTP inspection is not available without a special license.

Fields

- GTP Inspect Maps—Table that lists the defined GTP inspect maps.
- Add—Configures a new GTP inspect map. To edit a GTP inspect map, choose the GTP entry in the GTP Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the GTP Inspect Maps table.
- Security Level—Security level low only.
  - Do not Permit Errors
  - Maximum Number of Tunnels: 500
  - GSN timeout: 00:30:00
  - Pdp-Context timeout: 00:30:00
  - Request timeout: 00:01:00
  - Signaling timeout: 00:30:00.
  - Tunnel timeout: 01:00:00.
  - T3-response timeout: 00:00:20.
  - Drop and log unknown message IDs.
- IMSI Prefix Filtering—Opens the IMSI Prefix Filtering dialog box to configure IMSI prefix filters.
- Customize—Opens the Add/Edit GTP Policy Map dialog box for additional settings.
- Default Level—Sets the security level back to the default.

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

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
•	•	•	•	—

Add/Edit GTP Policy Map (Security Level)

The Add/Edit GTP Policy Map pane lets you configure the security level and additional settings for GTP application inspection maps.

Fields

- Name—When adding a GTP map, enter the name of the GTP map. When editing a GTP map, the name of the previously configured GTP map is shown.
- Description—Enter the description of the GTP map, up to 200 characters in length.
- Security Level—Security level low only.

  Do not Permit Errors

  Maximum Number of Tunnels: 500

  GSN timeout: 00:30:00

  Pdp-Context timeout: 00:30:00

  Request timeout: 00:01:00

  Signaling timeout: 00:30:00.

  Tunnel timeout: 01:00:00.

  T3-response timeout: 00:00:20.

  Drop and log unknown message IDs.

  – IMSI Prefix Filtering—Opens the IMSI Prefix Filtering dialog box to configure IMSI prefix filters.
  – Default Level—Sets the security level back to the default.
- Details—Shows the Parameters, IMSI Prefix Filtering, and Inspections tabs to configure additional settings.

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 Policy Map (Details)

The Add/Edit GTP Policy Map pane lets you configure the security level and additional settings for GTP application inspection maps.

Fields

- Name—When adding a GTP map, enter the name of the GTP map. When editing a GTP map, the name of the previously configured GTP map is shown.
- Description—Enter the description of the GTP map, up to 200 characters in length.
- Security Level—Shows the security level and IMSI prefix filtering settings to configure.
- Permit Parameters—Tab that lets you configure the permit parameters for the GTP inspect map.
  - Object Groups to Add

    From object group—Specify an object group or use the browse button to open the Add Network Object Group dialog box.

    To object group—Specify an object group or use the browse button to open the Add Network Object Group dialog box.
  - Add—Add the specified country code and network code to the IMSI Prefix table.
  - Delete—Deletes the specified country code and network code from the IMSI Prefix table.
  - Permit Errors—Lets any packets that are invalid or that encountered an error during inspection to be sent through the 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

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

The H.323 pane lets you view previously configured H.323 application inspection maps. An H.323 map lets you change the default configuration values used for H.323 application inspection.

H.323 inspection supports RAS, H.225, and H.245, and its functionality translates all embedded IP addresses and ports. It performs state tracking and filtering and can do a cascade of inspect function activation. H.323 inspection supports phone number filtering, dynamic T.120 control, H.245 tunneling control, HSI groups, protocol state tracking, H.323 call duration enforcement, and audio/video control.

Fields

- H.323 Inspect Maps—Table that lists the defined H.323 inspect maps.
- Add—Configures a new H.323 inspect map. To edit an H.323 inspect map, choose the H.323 entry in the H.323 Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the H.323 Inspect Maps table.
- Security Level—Select the security level (low, medium, or high).
  - Low—Default.

    State Checking h225 Disabled

    State Checking ras Disabled

    Call Party Number Disabled

    Call duration Limit Disabled

    RTP conformance not enforced
  - Medium

    State Checking h225 Enabled

    State Checking ras Enabled

    Call Party Number Disabled

    Call duration Limit Disabled

    RTP conformance enforced

    Limit payload to audio or video, based on the signaling exchange: no
  - High

    State Checking h225 Enabled

    State Checking ras Enabled

Call Party Number Enabled

Call duration Limit 1:00:00

RTP conformance enforced

Limit payload to audio or video, based on the signaling exchange: yes

- Phone Number Filtering—Opens the Phone Number Filtering dialog box to configure phone number filters.
- Customize—Opens the Add/Edit H.323 Policy Map dialog box for additional settings.
- Default Level—Sets the security level back to the default level of Medium.

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

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
•	•	•	•	—

Add/Edit H.323 Policy Map (Security Level)

The Add/Edit H.323 Policy Map pane lets you configure the security level and additional settings for H.323 application inspection maps.

Fields

- Name—When adding an H.323 map, enter the name of the H.323 map. When editing an H.323 map, the name of the previously configured H.323 map is shown.
- Description—Enter the description of the H323 map, up to 200 characters in length.
- Security Level—Select the security level (low, medium, or high).
  - Low—Default.

    State Checking h225 Disabled

    State Checking ras Disabled

    Call Party Number Disabled

    Call duration Limit Disabled

    RTP conformance not enforced
  - Medium

    State Checking h225 Enabled

    State Checking ras Enabled

    Call Party Number Disabled

    Call duration Limit Disabled

    RTP conformance enforced

    Limit payload to audio or video, based on the signaling exchange: no
  - High

    State Checking h225 Enabled

    State Checking ras Enabled

    Call Party Number Enabled

    Call duration Limit 1:00:00

    RTP conformance enforced

    Limit payload to audio or video, based on the signaling exchange: yes
  - Phone Number Filtering—Opens the Phone Number Filtering dialog box which lets you configure the settings for a phone number filter.
  - Default Level—Sets the security level back to the default.
- Details—Shows the State Checking, Call Attributes, Tunneling and Protocol Conformance, HSI Group Parameters, and Inspections tabs to configure additional settings.

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 Policy Map (Details)

The Add/Edit H.323 Policy Map pane lets you configure the security level and additional settings for H.323 application inspection maps.

Fields

- Name—When adding an H.323 map, enter the name of the H.323 map. When editing an H.323 map, the name of the previously configured H.323 map is shown.
- Description—Enter the description of the H.323 map, up to 200 characters in length.
- Security Level—Shows the security level and phone number filtering settings to configure.
- State Checking—Tab that lets you configure state checking parameters for the H.323 inspect map.
  - Check state transition of H.225 messages—Enforces H.323 state checking on H.225 messages.
  - Check state transition of RAS messages—Enforces H.323 state checking on RAS messages.
  - Check RFC messages and open pinholes for call signal addresses in RFQ messages

Note You can enable call setup between H.323 endpoints when the Gatekeeper is inside the network. The security appliance includes options to open pinholes for calls based on the RegistrationRequest/RegistrationConfirm (RRQ/RCF) messages. Because these RRQ/RCF messages are sent to and from the Gatekeeper, the calling endpoint's IP address is unknown and the security appliance opens a pinhole through source IP address/port 0/0. By default, this option is disabled. You can enable this option by setting the option in the H.323 Inspect Map.

- Call Attributes—Tab that lets you configure call attributes parameters for the H.323 inspect map.
  - Enforce call duration limit—Enforces the absolute limit on a call.

    Call Duration Limit—Time limit for the call (hh:mm:ss).
  - Enforce presence of calling and called party numbers—Enforces sending call party numbers during call setup.
- Tunneling and Protocol Conformance—Tab that lets you configure tunneling and protocol conformance parameters for the H.323 inspect map.
  - Check for H.245 tunneling—Allows H.245 tunneling.

    Action—Drop connection or log.
  - Check RTP packets for protocol conformance—Checks RTP/RTCP packets on the pinholes for protocol conformance.

    Limit payload to audio or video, based on the signaling exchange—Enforces the payload type to be audio or video based on the signaling exchange.

- HSI Group Parameters—Tab that lets you configure an HSI group.
  - HSI Group ID—Shows the HSI Group ID.
  - IP Address—Shows the HSI Group IP address.
  - Endpoints—Shows the HSI Group endpoints.
  - Add—Opens the Add HSI Group dialog box to add an HSI group.
  - Edit—Opens the Edit HSI Group dialog box to edit an HSI group.
  - Delete—Deletes an HSI group.
- Inspections—Tab that shows you the H.323 inspection configuration and lets you add or edit.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the H.323 inspection.
  - Value—Shows the value to match in the H.323 inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add H.323 Inspect dialog box to add an H.323 inspection.
  - Edit—Opens the Edit H.323 Inspect dialog box to edit an H.323 inspection.
  - Delete—Deletes an H.323 inspection.
  - Move Up—Moves an inspection up in the list.
  - Move Down—Moves an inspection down in the list.

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

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

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

The HTTP pane lets you view previously configured HTTP application inspection maps. An HTTP map lets you change the default configuration values used for HTTP application inspection.

HTTP application inspection scans HTTP headers and body, and performs various checks on the data. These checks prevent various HTTP constructs, content types, and tunneling and messaging protocols from traversing the security appliance.

HTTP application inspection can block tunneled applications and non-ASCII characters in HTTP requests and responses, preventing malicious content from reaching the web server. Size limiting of various elements in HTTP request and response headers, URL blocking, and HTTP server header type spoofing are also supported.

Fields

- HTTP Inspect Maps—Table that lists the defined HTTP inspect maps.
- Add—Configures a new HTTP inspect map. To edit an HTTP inspect map, choose the HTTP entry in the HTTP Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the HTTP Inspect Maps table.
- Security Level—Select the security level (low, medium, or high).
  - Low—Default.

    Protocol violation action: Drop connection

    Drop connections for unsafe methods: Disabled

    Drop connections for requests with non-ASCII headers: Disabled

    URI filtering: Not configured

    Advanced inspections: Not configured
  - Medium

    Protocol violation action: Drop connection

    Drop connections for unsafe methods: Allow only GET, HEAD, and POST

    Drop connections for requests with non-ASCII headers: Disabled

URI filtering: Not configured

Advanced inspections: Not configured

- High

  Protocol violation action: Drop connection and log

  Drop connections for unsafe methods: Allow only GET and HEAD.

  Drop connections for requests with non-ASCII headers: Enabled

  URI filtering: Not configured

  Advanced inspections: Not configured
- URI Filtering—Opens the URI Filtering dialog box to configure URI filters.
- Customize—Opens the Edit HTTP Policy Map dialog box for additional settings.
- Default Level—Sets the security level back to the default level of Medium.

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

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
•	•	•	•	—

Add/Edit HTTP Policy Map (Security Level)

The Add/Edit HTTP Policy Map pane lets you configure the security level and additional settings for HTTP application inspection maps.

Fields

- Name—When adding an HTTP map, enter the name of the HTTP map. When editing an HTTP map, the name of the previously configured HTTP map is shown.
- Description—Enter the description of the HTTP map, up to 200 characters in length.
- Security Level—Select the security level (low, medium, or high).
  - Low—Default.

    Protocol violation action: Drop connection

    Drop connections for unsafe methods: Disabled

    Drop connections for requests with non-ASCII headers: Disabled

    URI filtering: Not configured

    Advanced inspections: Not configured
  - Medium

    Protocol violation action: Drop connection

    Drop connections for unsafe methods: Allow only GET, HEAD, and POST

    Drop connections for requests with non-ASCII headers: Disabled

    URI filtering: Not configured

    Advanced inspections: Not configured
  - High

    Protocol violation action: Drop connection and log

    Drop connections for unsafe methods: Allow only GET and HEAD.

    Drop connections for requests with non-ASCII headers: Enabled

    URI filtering: Not configured

    Advanced inspections: Not configured
  - URI Filtering—Opens the URI Filtering dialog box which lets you configure the settings for an URI filter.
  - Default Level—Sets the security level back to the default.
- Details—Shows the Parameters and Inspections tabs to configure additional settings.

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 Policy Map (Details)

The Add/Edit HTTP Policy Map pane lets you configure the security level and additional settings for HTTP application inspection maps.

Fields

- Name—When adding an HTTP map, enter the name of the HTTP map. When editing an HTTP map, the name of the previously configured HTTP map is shown.
- Description—Enter the description of the HTTP map, up to 200 characters in length.
- Security Level—Shows the security level and URI filtering settings to configure.
- Parameters—Tab that lets you configure the parameters for the HTTP inspect map.
  - Check for protocol violations—Checks for HTTP protocol violations.

    Action—Drop Connection, Reset, Log.

    Log—Enable or disable.
  - Spoof server string—Replaces the server HTTP header value with the specified string.

    Spoof String—Enter a string to substitute for the server header field. Maximum is 82 characters.
  - Body Match Maximum—The maximum number of characters in the body of an HTTP message that should be searched in a body match. Default is 200 bytes. A large number will have a significant impact on performance.
- Inspections—Tab that shows you the HTTP inspection configuration and lets you add or edit.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the HTTP inspection.
  - Value—Shows the value to match in the HTTP inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add HTTP Inspect dialog box to add an HTTP inspection.
  - Edit—Opens the Edit HTTP Inspect dialog box to edit an HTTP inspection.
  - Delete—Deletes an HTTP inspection.
  - Move Up—Moves an inspection up in the list.
  - Move Down—Moves an inspection down in the list.

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

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

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.
- Add—Configures a new IM inspect map.
- Edit—Edits the selected IM entry in the IM Inspect Maps table.
- Delete—Deletes the inspect map selected in the IM Inspect Maps 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 Instant Messaging (IM) Policy Map

The Add/Edit Instant Messaging (IM) Policy Map pane lets you configure the security level and additional settings for IM application inspection maps.

Fields

- Name—When adding an IM map, enter the name of the IM map. When editing an IM map, the name of the previously configured IM map is shown.
- 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

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

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

- IPSec Pass Through Inspect Maps—Table that lists the defined IPSec Pass Through inspect maps.
- Add—Configures a new IPSec Pass Through inspect map. To edit an IPSec Pass Through inspect map, select the IPSec Pass Through entry in the IPSec Pass Through Inspect Maps table and click Customize.
- Delete—Deletes the inspect map selected in the IPSec Pass Through Inspect Maps table.
- Security Level—Select the security level (high or low).

- Low—Default.

  Maximum ESP flows per client: Unlimited.

  ESP idle timeout: 00:10:00.

  Maximum AH flows per client: Unlimited.

  AH idle timeout: 00:10:00.

- High

  Maximum ESP flows per client:10.

  ESP idle timeout: 00:00:30.

  Maximum AH flows per client: 10.

  AH idle timeout: 00:00:30.

- Customize—Opens the Add/Edit IPSec Pass Thru Policy Map dialog box for additional settings.
- Default Level—Sets the security level back to the default level of Low.

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 IPSec Pass Thru Policy Map (Security Level)

The Add/Edit IPSec Pass Thru Policy Map pane lets you configure the security level and additional settings for IPSec Pass Thru application inspection maps.

Fields

- Name—When adding an IPSec Pass Thru map, enter the name of the IPSec Pass Thru map. When editing an IPSec Pass Thru map, the name of the previously configured IPSec Pass Thru map is shown.
- Security Level—Select the security level (high or low).
  - Low—Default.

    Maximum ESP flows per client: Unlimited.

    ESP idle timeout: 00:10:00.

    Maximum AH flows per client: Unlimited.

    AH idle timeout: 00:10:00.

  - High

    Maximum ESP flows per client:10.

    ESP idle timeout: 00:00:30.

    Maximum AH flows per client: 10.

AH idle timeout: 00:00:30.

  - Default Level—Sets the security level back to the default level of Low.

- Details—Shows additional parameter settings to configure.

Mode

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

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

Add/Edit IPSec Pass Thru Policy Map (Details)

The Add/Edit IPSec Pass Thru Policy Map pane lets you configure the security level and additional settings for IPSec Pass Thru application inspection maps.

Fields

- Name—When adding an IPSec Pass Thru map, enter the name of the IPSec Pass Thru map. When editing an IPSec Pass Thru map, the name of the previously configured IPSec Pass Thru map is shown.
- Description—Enter the description of the IPSec Pass Through map, up to 200 characters in length.
- Security Level—Shows the security level settings to configure.
- Parameters—Configures ESP and AH parameter settings.
  - Limit ESP flows per client—Limits ESP flows per client.

    Maximum—Specify maximum limit.
  - Apply ESP idle timeout—Applies ESP idle timeout.

    Timeout—Specify timeout.
  - Limit AH flows per client—Limits AH flows per client.

    Maximum—Specify maximum limit.
  - Apply AH idle timeout—Applies AH idle timeout.

    Timeout—Specify timeout.

Modes

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

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

IP Options Inspect Map

The IP Options Inspect Maps pane lets you view previously configured IP Options inspection maps. An IP Options inspection map lets you change the default configuration values used for IP Option inspection.

You can configure IP Options inspection to control which IP packets with specific IP options are allowed through the security appliance. Configuring this inspection instructs the security appliance to allow a packet to pass or to clear the specified IP options and then allow the packet to pass.

In particular, you can control whether the security appliance drops, clears, or passes packets containing the Router Alert (RTRALT) option. Dropping RSVP packets containing the Router Alert option can cause problems in VoIP implementations. Therefore, you can create IP Options inspection maps to pass packets containing the RTRALT option.

Fields

IP Options Inspect Maps—Table that lists the defined IP Options inspect maps.

Add—Configures a new IP Options inspect map.

Edit—Edits an existing IP Options inspect map. To edit an IP Options inspect map, choose the entry in the table and click Edit.

Delete—Deletes the inspect map selected in the IP-Options Inspect Maps table.

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 Options Inspect Map

The Add/Edit IP Options Inspect Map lets you configure the settings for IP Options inspection maps.

Fields

- Name—When adding an IP Options inspection map, enter the name of the map. When editing a map, the name of the previously configured map is shown.
- Description—Enter the description of the IP Options inspection map, up to 200 characters in length.
- Parameters—Select which IP options you want to pass through the security appliance or clear and then pass through the security appliance:
  - Allow packets with the End of Options List (EOOL) option

    This option, which contains just a single zero byte, appears at the end of all options to mark the end of a list of options. This might not coincide with the end of the header according to the header length.
  - Allow packets with the No Operation (NOP) option

The Options field in the IP header can contain zero, one, or more options, which makes the total length of the field variable. However, the IP header must be a multiple of 32 bits. If the number of bits of all options is not a multiple of 32 bits, the NOP option is used as “internal padding” to align the options on a 32-bit boundary.

– Allow packets with the Router Alert (RTRALT) option

This option notifies transit routers to inspect the contents of the packet even when the packet is not destined for that router. This inspection is valuable when implementing RSVP and similar protocols require relatively complex processing from the routers along the packets delivery path.

– Clear the option value from the packets

When an option is checked, the **Clear the option value from the packets** check box becomes available for that option. Select the **Clear the option value from the packets** check box to clear the option from the packet before allowing the packet through the 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
•	•	•	•	—

MGCP Inspect Map

The MGCP pane lets you view previously configured MGCP application inspection maps. An MGCP map lets you change the default configuration values used for MGCP application inspection. You can use an MGCP map to manage connections between VoIP devices and MGCP call agents.

Fields

- MGCP Inspect Maps—Table that lists the defined MGCP inspect maps.
- Add—Configures a new MGCP inspect map.
- Edit—Edits the selected MGCP entry in the MGCP Inspect Maps table.
- Delete—Deletes the inspect map selected in the MGCP Inspect Maps table.

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

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
•	•	•	•	—

Add/Edit MGCP Policy Map

The Add/Edit MGCP Policy Map pane lets you configure the command queue, gateway, and call agent settings for MGCP application inspection maps.

Fields

- Name—When adding an MGCP map, enter the name of the MGCP map. When editing an MGCP map, the name of the previously configured MGCP map is shown.
- Description—Enter the description of the MGCP map, up to 200 characters in length.
- Command Queue—Tab that lets you specify the permitted queue size for MGCP commands.
  - Command Queue Size—Specifies the maximum number of commands to queue. The valid range is from 1 to 2147483647.
- Gateways and Call Agents—Tab that lets you configure groups of gateways and call agents for this map.

- Group ID—Identifies the ID of the call agent group. A call agent group associates one or more call agents with one or more MGCP media gateways. The gateway IP address can only be associated with one group ID. You cannot use the same gateway with different group IDs. The valid range is from 0 to 2147483647Criterion—Shows the criterion of the inspection.
- Gateways—Identifies the IP address of the media gateway that is controlled by the associated call agent. A media gateway is typically a network element that provides conversion between the audio signals carried on telephone circuits and data packets carried over the Internet or over other packet networks. Normally, a gateway sends commands to the default MGCP port for call agents, 2727.
- Call Agents—Identifies the IP address of a call agent that controls the MGCP media gateways in the call agent group. Normally, a call agent sends commands to the default MGCP port for gateways, 2427.
- Add—Displays the Add MGCP Group dialog box, which you can use to define a new MGCP group of gateways and call agents.
- Edit—Displays the Edit MGCP dialog box, which you can use to modify the MGCP group selected in the Gateways and Call Agents table.
- Delete—Deletes the MGCP group selected in the Gateways and Call Agents table.

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

The Add/Edit MGCP Group dialog box lets you define the configuration of an MGCP group that will be used when MGCP application inspection is enabled.

Fields

- Group ID—Specifies the ID of the call agent group. A call agent group associates one or more call agents with one or more MGCP media gateways. The valid range is from 0 to 2147483647.
  - Gateway to Be Added—Specifies the IP address of the media gateway that is controlled by the associated call agent. A media gateway is typically a network element that provides conversion between the audio signals carried on telephone circuits and data packets carried over the Internet or over other packet networks. Normally, a gateway sends commands to the default MGCP port for call agents, 2727.
  - Add—Adds the specified IP address to the IP address table.
  - Delete—Deletes the selected IP address from the IP address table.
  - IP Address—Lists the IP addresses of the gateways in the call agent group.
- Call Agents

  - Call Agent to Be Added—Specifies the IP address of a call agent that controls the MGCP media gateways in the call agent group. Normally, a call agent sends commands to the default MGCP port for gateways, 2427.
  - Add—Adds the specified IP address to the IP address table.
  - Delete—Deletes the selected IP address from the IP address table.
  - IP Address—Lists the IP addresses of the call agents in the call agent group.

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

The NetBIOS pane lets you view previously configured NetBIOS application inspection maps. A NetBIOS map lets you change the default configuration values used for NetBIOS application inspection.

NetBIOS application inspection performs NAT for the embedded IP address in the NetBIOS name service packets and NetBIOS datagram services packets. It also enforces protocol conformance, checking the various count and length fields for consistency.

Fields

- NetBIOS Inspect Maps—Table that lists the defined NetBIOS inspect maps.
- Add—Configures a new NetBIOS inspect map.
- Edit—Edits the selected NetBIOS entry in the NetBIOS Inspect Maps table.
- Delete—Deletes the inspect map selected in the NetBIOS Inspect Maps table.

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 NetBIOS Policy Map

The Add/Edit NetBIOS Policy Map pane lets you configure the protocol violation settings for NetBIOS application inspection maps.

Fields

- Name—When adding a NetBIOS map, enter the name of the NetBIOS map. When editing an NetBIOS map, the name of the previously configured NetBIOS map is shown.
- Description—Enter the description of the NetBIOS map, up to 200 characters in length.
- Check for protocol violations—Checks for protocol violations and executes specified action.
  - Action—Drop packet or log.
  - Log—Enable or disable.

Modes

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

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

RTSP Inspect Map

The RTSP pane lets you view previously configured RTSP application inspection maps. An RTSP map lets you change the default configuration values used for RTSP application inspection. You can use an RTSP map to protect RTSP traffic.

Fields

- RTSP Inspect Maps—Table that lists the defined RTSP inspect maps.
- Add—Configures a new RTSP inspect map.
- Edit—Edits the selected RTSP entry in the RTSP Inspect Maps table.
- Delete—Deletes the inspect map selected in the RTSP Inspect Maps table.

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 RTSP Policy Map

The Add/Edit RTSP Policy Map pane lets you configure the parameters and inspections settings for RTSP application inspection maps.

Fields

- Name—When adding an RTSP map, enter the name of the RTSP map. When editing an RTSP map, the name of the previously configured RTSP map is shown.
- Description—Enter the description of the RTSP map, up to 200 characters in length.
- Parameters—Tab that lets you restrict usage on reserved ports during media port negotiation, and lets you set the URL length limit.
  - Enforce Reserve Port Protection—Lets you restrict the use of reserved ports during media port negotiation.
  - Maximum URL Length—Specifies the maximum length of the URL allowed in the message. Maximum value is 6000.
- Inspections—Tab that shows you the RTSP inspection configuration and lets you add or edit.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the RTSP inspection.
  - Value—Shows the value to match in the RTSP inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add RTSP Inspect dialog box to add a RTSP inspection.
  - Edit—Opens the Edit RTSP Inspect dialog box to edit a RTSP inspection.
  - Delete—Deletes a RTSP inspection.
  - Move Up—Moves an inspection up in the list.
  - Move Down—Moves an inspection down in the list.

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 RTSP Inspect

The Add/Edit RTSP Inspect dialog box lets you define the match criterion, values, and actions for the RTSP 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 RTSP traffic to match.
  - URL Filter—Match URL filtering.

- Request Method—Match an RTSP request method.
- URL Filter Criterion Values—Specifies to match URL filtering. Applies the regular expression match.
  - Regular Expression—Lists the defined regular expressions to match.
  - Manage—Opens the Manage Regular Expressions dialog box, which lets you configure regular expressions.
  - Regular Expression Class—Lists the defined regular expression classes to match.
  - Manage—Opens the Manage Regular Expression Class dialog box, which lets you configure regular expression class maps.
- URL Filter Actions—Primary action and log settings.
  - Action—Drop connection or log.
  - Log—Enable or disable.
- Request Method Criterion Values—Specifies to match an RTSP request method.
  - Request Method—Specifies a request method: announce, describe, get_parameter, options, pause, play, record, redirect, setup, set_parameters, teardown.
- Request Method Actions—Primary action settings.
  - Action—Limit rate (pps).

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

The SCCP (Skinny) pane lets you view previously configured SCCP (Skinny) application inspection maps. An SCCP (Skinny) map lets you change the default configuration values used for SCCP (Skinny) application inspection.

Skinny application inspection performs translation of embedded IP address and port numbers within the packet data, and dynamic opening of pinholes. It also performs additional protocol conformance checks and basic state tracking.

Fields

- SCCP (Skinny) Inspect Maps—Table that lists the defined SCCP (Skinny) inspect maps.
- Add—Configures a new SCCP (Skinny) inspect map. To edit an SCCP (Skinny) inspect map, choose the SCCP (Skinny) entry in the SCCP (Skinny) Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the SCCP (Skinny) Inspect Maps table.
- Security Level—Select the security level (high or low).
  - Low—Default.

Registration: Not enforced.

Maximum message ID: 0x181.

Minimum prefix length: 4

Media timeout: 00:05:00

Signaling timeout: 01:00:00.

RTP conformance: Not enforced.

- Medium

  Registration: Not enforced.

  Maximum message ID: 0x141.

  Minimum prefix length: 4.

  Media timeout: 00:01:00.

  Signaling timeout: 00:05:00.

  RTP conformance: Enforced.

  Limit payload to audio or video, based on the signaling exchange: No.

- High

  Registration: Enforced.

  Maximum message ID: 0x141.

  Minimum prefix length: 4.

  Maximum prefix length: 65536.

  Media timeout: 00:01:00.

  Signaling timeout: 00:05:00.

  RTP conformance: Enforced.

  Limit payload to audio or video, based on the signaling exchange: Yes.

- Message ID Filtering—Opens the Messaging ID Filtering dialog box for configuring message ID filters.
- Customize—Opens the Add/Edit SCCP (Skinny) Policy Map dialog box for additional settings.
- Default Level—Sets the security level back to the default level of Low.

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

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
•	•	•	•	—

Add/Edit SCCP (Skinny) Policy Map (Security Level)

The Add/Edit SCCP (Skinny) Policy Map pane lets you configure the security level and additional settings for SCCP (Skinny) application inspection maps.

Fields

- Name—When adding an SCCP (Skinny) map, enter the name of the SCCP (Skinny) map. When editing an SCCP (Skinny) map, the name of the previously configured SCCP (Skinny) map is shown.
- Description—Enter the description of the SCCP (Skinny) map, up to 200 characters in length.
- Security Level—Select the security level (high or low).
  - Low—Default.

    Registration: Not enforced.

    Maximum message ID: 0x181.

    Minimum prefix length: 4

    Media timeout: 00:05:00

    Signaling timeout: 01:00:00.

    RTP conformance: Not enforced.
  - Medium

    Registration: Not enforced.

    Maximum message ID: 0x141.

Minimum prefix length: 4.

Media timeout: 00:01:00.

Signaling timeout: 00:05:00.

RTP conformance: Enforced.

Limit payload to audio or video, based on the signaling exchange: No.

- High

  Registration: Enforced.

  Maximum message ID: 0x141.

  Minimum prefix length: 4.

  Maximum prefix length: 65536.

  Media timeout: 00:01:00.

  Signaling timeout: 00:05:00.

  RTP conformance: Enforced.

  Limit payload to audio or video, based on the signaling exchange: Yes.

- Message ID Filtering—Opens the Messaging ID Filtering dialog box for configuring message ID filters.
- Default Level—Sets the security level back to the default.

- Details—Shows additional parameter, RTP conformance, and message ID filtering settings to configure.

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 SCCP (Skinny) Policy Map (Details)

The Add/Edit SCCP (Skinny) Policy Map pane lets you configure the security level and additional settings for SCCP (Skinny) application inspection maps.

Fields

- Name—When adding an SCCP (Skinny) map, enter the name of the SCCP (Skinny) map. When editing an SCCP (Skinny) map, the name of the previously configured SCCP (Skinny) map is shown.
- Description—Enter the description of the DNS map, up to 200 characters in length.
- Security Level—Shows the security level and message ID filtering settings to configure.
- Parameters—Tab that lets you configure the parameter settings for SCCP (Skinny).
  - Enforce endpoint registration—Enforce that Skinny endpoints are registered before placing or receiving calls.

  Maximum Message ID—Specify value of maximum SCCP message ID allowed.

  - SCCP Prefix Length—Specifies prefix length value in Skinny messages.

    Minimum Prefix Length—Specify minimum value of SCCP prefix length allowed.

    Maximum Prefix Length—Specify maximum value of SCCP prefix length allowed.

  - Media Timeout—Specify timeout value for media connections.
  - Signaling Timeout—Specify timeout value for signaling connections.

- RTP Conformance—Tab that lets you configure the RTP conformance settings for SCCP (Skinny).
  - Check RTP packets for protocol conformance—Checks RTP/RTCP packets flowing on the pinholes for protocol conformance.

    Limit payload to audio or video, based on the signaling exchange—Enforces the payload type to be audio/video based on the signaling exchange.

- Message ID Filtering—Tab that lets you configure the message ID filtering settings for SCCP (Skinny).
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the inspection.
  - Value—Shows the value to match in the inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add Message ID Filtering dialog box to add a message ID filter.
  - Edit—Opens the Edit Message ID Filtering dialog box to edit a message ID filter.
  - Delete—Deletes a message ID filter.
  - Move Up—Moves an entry up in the list.
  - Move Down—Moves an entry down in the list.

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

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

The SIP pane lets you view previously configured SIP application inspection maps. A SIP map lets you change the default configuration values used for SIP application inspection.

SIP is a widely used protocol for Internet conferencing, telephony, presence, events notification, and instant messaging. Partially because of its text-based nature and partially because of its flexibility, SIP networks are subject to a large number of security threats.

SIP application inspection provides address translation in message header and body, dynamic opening of ports and basic sanity checks. It also supports application security and protocol conformance, which enforce the sanity of the SIP messages, as well as detect SIP-based attacks.

Fields

- SIP Inspect Maps—Table that lists the defined SIP inspect maps.
- Add—Configures a new SIP inspect map. To edit a SIP inspect map, choose the SIP entry in the SIP Inspect Maps table and click **Customize**.
- Delete—Deletes the inspect map selected in the SIP Inspect Maps table.
- Security Level—Select the security level (high or low).
  - Low—Default.

    SIP instant messaging (IM) extensions: Enabled.

    Non-SIP traffic on SIP port: Permitted.

    Hide server's and endpoint's IP addresses: Disabled.

    Mask software version and non-SIP URIs: Disabled.

    Ensure that the number of hops to destination is greater than 0: Enabled.

    RTP conformance: Not enforced.

SIP conformance: Do not perform state checking and header validation.

- Medium

  SIP instant messaging (IM) extensions: Enabled.

  Non-SIP traffic on SIP port: Permitted.

  Hide server's and endpoint's IP addresses: Disabled.

  Mask software version and non-SIP URIs: Disabled.

  Ensure that the number of hops to destination is greater than 0: Enabled.

  RTP conformance: Enforced.

  Limit payload to audio or video, based on the signaling exchange: No

  SIP conformance: Drop packets that fail state checking.

- High

  SIP instant messaging (IM) extensions: Enabled.

  Non-SIP traffic on SIP port: Denied.

  Hide server's and endpoint's IP addresses: Disabled.

  Mask software version and non-SIP URIs: Enabled.

  Ensure that the number of hops to destination is greater than 0: Enabled.

  RTP conformance: Enforced.

  Limit payload to audio or video, based on the signaling exchange: Yes

  SIP conformance: Drop packets that fail state checking and packets that fail header validation.

- Customize—Opens the Add/Edit SIP Policy Map dialog box for additional settings.
- Default Level—Sets the security level back to the default level of Low.

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 Policy Map (Security Level)

The Add/Edit SIP Policy Map pane lets you configure the security level and additional settings for SIP application inspection maps.

Fields

- Name—When adding a SIP, enter the name of the SIP map. When editing a SIP map, the name of the previously configured SIP map is shown.
- Description—Enter the description of the SIP map, up to 200 characters in length.
- Security Level—Select the security level (high or low).

  - Low—Default.

    SIP instant messaging (IM) extensions: Enabled.

    Non-SIP traffic on SIP port: Permitted.

    Hide server's and endpoint's IP addresses: Disabled.

    Mask software version and non-SIP URIs: Disabled.

    Ensure that the number of hops to destination is greater than 0: Enabled.

    RTP conformance: Not enforced.

    SIP conformance: Do not perform state checking and header validation.

  - Medium

    SIP instant messaging (IM) extensions: Enabled.

    Non-SIP traffic on SIP port: Permitted.

    Hide server's and endpoint's IP addresses: Disabled.

    Mask software version and non-SIP URIs: Disabled.

    Ensure that the number of hops to destination is greater than 0: Enabled.

    RTP conformance: Enforced.

    Limit payload to audio or video, based on the signaling exchange: No

    SIP conformance: Drop packets that fail state checking.

  - High

    SIP instant messaging (IM) extensions: Enabled.

    Non-SIP traffic on SIP port: Denied.

    Hide server's and endpoint's IP addresses: Disabled.

    Mask software version and non-SIP URIs: Enabled.

    Ensure that the number of hops to destination is greater than 0: Enabled.

    RTP conformance: Enforced.

    Limit payload to audio or video, based on the signaling exchange: Yes

    SIP conformance: Drop packets that fail state checking and packets that fail header validation.

  - Default Level—Sets the security level back to the default.

- Details—Shows additional filtering, IP address privacy, hop count, RTP conformance, SIP conformance, field masking, and inspections settings to configure.

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 Policy Map (Details)

The Add/Edit SIP Policy Map pane lets you configure the security level and additional settings for SIP application inspection maps.

Fields

- Name—When adding a SIP, enter the name of the SIP map. When editing a SIP map, the name of the previously configured SIP map is shown.
- Description—Enter the description of the SIP map, up to 200 characters in length.
- Security Level—Shows the security level settings to configure
- Filtering—Tab that lets you configure the filtering settings for SIP.
  - Enable SIP instant messaging (IM) extensions—Enables Instant Messaging extensions. Default is enabled.
  - Permit non-SIP traffic on SIP port—Permits non-SIP traffic on SIP port. Permitted by default.
- IP Address Privacy—Tab that lets you configure the IP address privacy settings for SIP.
  - Hide server's and endpoint's IP addresses—Enables IP address privacy. Disabled by default.
- Hop Count—Tab that lets you configure the hop count settings for SIP.
  - Ensure that number of hops to destination is greater than 0—Enables check for the value of Max-Forwards header is zero.

    Action—Drop packet, Drop Connection, Reset, Log.

    Log—Enable or Disable.
- RTP Conformance—Tab that lets you configure the RTP conformance settings for SIP.
  - Check RTP packets for protocol conformance—Checks RTP/RTCP packets flowing on the pinholes for protocol conformance.

    Limit payload to audio or video, based on the signaling exchange—Enforces payload type to be audio/video based on the signaling exchange.
- SIP Conformance—Tab that lets you configure the SIP conformance settings for SIP.
  - Enable state transition checking—Enables SIP state checking.

    Action—Drop packet, Drop Connection, Reset, Log.

    Log—Enable or Disable.
  - Enable strict validation of header fields—Enables validation of SIP header fields.

    Action—Drop packet, Drop Connection, Reset, Log.

    Log—Enable or Disable.
- Field Masking—Tab that lets you configure the field masking settings for SIP.
  - Inspect non-SIP URIs—Enables non-SIP URI inspection in Alert-Info and Call-Info headers.

    Action—Mask or Log.

    Log—Enable or Disable.
  - Inspect server's and endpoint's software version—Inspects SIP endpoint software version in User-Agent and Server headers.

    Action—Mask or Log.

    Log—Enable or Disable.

- Inspections—Tab that shows you the SIP inspection configuration and lets you add or edit.
  - Match Type—Shows the match type, which can be a positive or negative match.
  - Criterion—Shows the criterion of the SIP inspection.
  - Value—Shows the value to match in the SIP inspection.
  - Action—Shows the action if the match condition is met.
  - Log—Shows the log state.
  - Add—Opens the Add SIP Inspect dialog box to add a SIP inspection.
  - Edit—Opens the Edit SIP Inspect dialog box to edit a SIP inspection.
  - Delete—Deletes a SIP inspection.
  - Move Up—Moves an inspection up in the list.
  - Move Down—Moves an inspection down in the list.

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

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

The SNMP pane lets you view previously configured SNMP application inspection maps. An SNMP map lets you change the default configuration values used for SNMP application inspection.

Fields

- Map Name—Lists previously configured application inspection maps. Select a map and click **Edit** to view or change an existing map.
- Add—Configures a new SNMP inspect map.

- Edit—Edits the selected SNMP entry in the SNMP Inspect Maps table.
- Delete—Deletes the inspect map selected in the SNMP Inspect Maps table.

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

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
•	•	•	•	—

Encrypted Traffic Inspection

This section describes how to configure encrypted traffic inspection, and includes the following topics:

- TLS Proxy Wizard, page 25-133
- Phone Proxy, page 25-142
- CTL File, page 25-147
- TLS Proxy, page 25-150
- CTL Provider, page 25-152

TLS Proxy Wizard

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

For information on how to configure the TLS Proxy, see the following sections:

- Configure TLS Proxy Pane, page 25-134
- Adding a TLS Proxy Instance, page 25-135
- Add TLS Proxy Instance Wizard – Server Configuration, page 25-136
- Add TLS Proxy Instance Wizard – Client Configuration, page 25-137
- Add TLS Proxy Instance Wizard – Other Steps, page 25-139

Use the TLS Proxy to enable inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager. Additionally, configure the TLS Proxy on the security appliance to use the following Cisco Unified Communications features:

Table 25-2 TLS Proxy Applications and the Security Appliance

Application	TLS Client	TLS Server	Client Authentication	Security Appliance Server Role	Security Appliance Client Role
Mobile Advantage	CUMC	CUMA	No	Using the CUMA private key or certificate impersonation	Any static configured certificate
Presence Federation	CUP or MS LCS/OCS	CUP or MS LCS/OCS	Yes	Proxy certificate, self-signed or by internal CA	Using the CUP private key or certificate impersonation
IP Telephone (including Phone Proxy)	IP phone	CUCM	Yes	Proxy certificate, self-signed or by internal CA	Local dynamic certificate signed by the security appliance CA (might not need certificate for Phone Proxy application)

For the Mobility feature, the TLS client is a CUMA client and the TLS server is a CUMA server. The security appliance is between a CUMA client and a CUMA server. The TLS Proxy for CUMA allows the use of an imported PKCS-12 certificate for server proxy during the handshake with the client. CUMA clients are not required to present a certificate (no client authentication) during the handshake. In previous releases, the security appliance required the client to always present a valid certificate and it acted as a private certificate authority (CA) for the clients.

For the Presence Federation feature, the security appliance acts as a TLS Proxy between the Cisco Unified Presence and the foreign server. This allows the security appliance to proxy TLS messages on behalf of the server that initiates the TLS connection, and route the proxied TLS messages to the client. The security appliance stores certificate trustpoints for the server and the client, and presents these certificates on establishment of the TLS session.

The security appliance supports TLS Proxy for various voice applications. For the Phone Proxy feature, the TLS Proxy running on the security appliance has the following key features:

- The TLS Proxy is implemented on the security appliance to intercept the TLS signaling from IP phones.
- The TLS Proxy decrypts the packets, sends packets to the inspection engine for NAT rewrite and protocol conformance, optionally encrypts packets, and sends them to CUCM or sends them in clear text if the IP phone is configured to be in nonsecure mode on the CUCM.
- The TLS Proxy is a transparent proxy that works based on establishing trusted relationship between the TLS client, the proxy (the security appliance), and the TLS Server.

Configure TLS Proxy Pane

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

You can configure the TLS Proxy from the Configuration > Firewall > Advanced > Encrypted Traffic Inspection > TLS Proxy pane. For a detailed overview of the TLS Proxy, see TLS Proxy Wizard, page 25-133.

Configuring a TLS Proxy lets you use the TLS Proxy to enable inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and enable the security appliance for the Cisco Unified Communications features:

- TLS Proxy for the Cisco Unified Presence Server (CUPS), part of Presence Federation
- TLS Proxy for the Cisco Unified Mobility Advantage (CUMA) server, part of Mobile Advantage
- Phone Proxy

Fields

- TLS Proxy Name—Lists the TLS Proxy name.
- Server Proxy Certificate—Lists the trustpoint, which is either self-signed or enrolled with a certificate server.
- Local Dynamic Certificate Issuer—Lists the local certificate authority to issue client or server dynamic certificates.
- Client Proxy Certificate—Lists the proxy certificate for the TLS client. The security appliance uses the client proxy certificate to authenticate the TLS client during the handshake between the proxy and the TLS client. The certificate can be either self-signed, enrolled with a certificate authority, or issued by the third party.
- Add—Adds a TLS Proxy by launching the Add TLS Proxy Instance Wizard. See Adding a TLS Proxy Instance, page 25-135 for the steps to create a TLS Proxy instance.
- Edit—Edits a TLS Proxy. The fields in the Edit panel area identical to the fields displayed when you add a TLS Proxy instance. See Edit TLS Proxy Instance – Server Configuration, page 25-139 and Edit TLS Proxy Instance – Client Configuration, page 25-140.
- Delete—Deletes a TLS Proxy.

- Maximum Sessions—Lets you specify the maximum number of TLS Proxy sessions to support.
  - Specify the maximum number of TLS Proxy sessions that the ASA needs to support.
  - Maximum number of sessions—The minimum is 1. The maximum is dependent on the platform:

    Cisco ASA 5505 security appliance: 10

    Cisco ASA 5510 security appliance: 100

    Cisco ASA 5520 security appliance: 300

    Cisco ASA 5540 security appliance: 1000

    Cisco ASA 5550 security appliance: 2000

    Cisco ASA 5580 security appliance: 4000

Note The maximum number of sessions is global to all TLS proxy sessions.

Modes

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

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

Adding a TLS Proxy Instance

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Use the Add TLS Proxy Instance Wizard to add a TLS Proxy to enable inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and to support the Cisco Unified Communications features on the security appliance. For a detailed overview of the TLS Proxy used by these features, see TLS Proxy Wizard, page 25-133.

This wizard is available from the Configuration > Firewall > Advanced > Encrypted Traffic Inspection > TLS Proxy pane.

Step 1 Open the Configuration > Firewall > Advanced > Encrypted Traffic Inspection > TLS Proxy pane.

Step 2 To add a new TLS Proxy Instance, click **Add**.

The Add TLS Proxy Instance Wizard opens.

Step 3 In the TLS Proxy Name field, type the TLS Proxy name.

Step 4 Click **Next**.

The Add TLS Proxy Instance Wizard – Server Configuration dialog box opens. In this step of the wizard, configure the server proxy parameters for original TLS Server—the Cisco Unified Call Manager (CUCM) server, the Cisco Unified Presence Server (CUPS), or the Cisco Unified Mobility Advantage (CUMA) server. See Add TLS Proxy Instance Wizard – Server Configuration, page 25-136.

After configuring the server proxy parameters, the wizard guides you through configuring client proxy parameters (see Add TLS Proxy Instance Wizard – Client Configuration, page 25-137) and provides instructions on the steps to complete outside the ASDM to make the TLS Proxy fully functional (see Add TLS Proxy Instance Wizard – Other Steps, page 25-139).

Modes

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

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

Add TLS Proxy Instance Wizard – Server Configuration

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Use the Add TLS Proxy Instance Wizard to add a TLS Proxy to enable inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and to support the Cisco Unified Communications features on the security appliance. For a detailed overview of the TLS Proxy used by these features, see TLS Proxy Wizard, page 25-133.

The Add TLS Proxy Instance Wizard is available from the Configuration > Firewall > Advanced > Encrypted Traffic Inspection > TLS Proxy pane.

Step 1 Complete the first step of the Add TLS Proxy Instance Wizard. See Adding a TLS Proxy Instance, page 25-135.

The Add TLS Proxy Instance Wizard – Server Configuration dialog box opens.

Step 2 Specify the server proxy certificate by doing one of the following:

- To add a new certificate, click **Manage**. The Manage Identify Certificates dialog box opens.

  When the Phone Proxy is operating in a mixed-mode CUCM cluster, you must import the CUCM certificate by clicking **Add** in the Manage Identify Certificates dialog box. See Adding or Importing an Identity Certificate, page 34-9.

- To select an existing certificate, select one from the drop-down list.

  When you are configuring the TLS Proxy for the Phone Proxy, select the certificate that has a filename beginning with **_internal_PP_**. When you create the CTL file for the Phone Proxy, the security appliance, creates an internal trustpoint used by the Phone Proxy to sign the TFTP files. The trustpoint is named **_internal_PP_***ctl-instance_filename*.

The server proxy certificate is used to specify the trustpoint to present during the TLS handshake. The trustpoint can be self-signed or enrolled locally with the certificate service on the proxy. For example, for the Phone Proxy, the server proxy certificate is used by the Phone Proxy during the handshake with the IP phones.

Step 3 To install the TLS server certificate in the security appliance trust store, so that the security appliance can authenticate the TLS server during TLS handshake between the proxy and the TLS server, click **Install TLS Server's Certificate**.

The Manage CA Certificates dialog box opens. See Guidelines and Limitations, page 34-2. Click **Add** to open the Install Certificate dialog box. See Adding or Installing a CA Certificate, page 34-3.

When you are configuring the TLS Proxy for the Phone Proxy, click **Install TLS Server's Certificate** and install the Cisco Unified Call Manager (CUCM) certificate so that the proxy can authenticate the IP phones on behalf of the CUCM server.

Step 4 To require the security appliance to present a certificate and authenticate the TLS client during TLS handshake, check the Enable client authentication during TLS Proxy handshake check box.

When adding a TLS Proxy Instance for Mobile Advantage (the CUMC client and CUMA server), disable the check box when the client is incapable of sending a client certificate.

See TLS Proxy Wizard, page 25-133 to determine which TLS clients used by the Cisco Unified Communication features are capable of client authentication.

Step 5 Click **Next**.

The Add TLS Proxy Instance Wizard – Client Configuration dialog box opens. In this step of the wizard, configure the client proxy parameters for original TLS Client—the CUMC client for Mobile Advantage, CUP or MS LCS/OCS client for Presence Federation, or the IP phone for the Phone Proxy. See Add TLS Proxy Instance Wizard – Client Configuration, page 25-137.

After configuring the client proxy parameters, the wizard provides instructions on the steps to complete outside the ASDM to make the TLS Proxy fully functional (see Add TLS Proxy Instance Wizard – Other Steps, page 25-139).

Add TLS Proxy Instance Wizard – Client Configuration

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Use the Add TLS Proxy Instance Wizard to add a TLS Proxy to enable inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and to support the Cisco Unified Communications features on the security appliance. For a detailed overview of the TLS Proxy used by these features, see TLS Proxy Wizard, page 25-133.

This wizard is available from the Configuration > Firewall > Advanced > Encrypted Traffic Inspection > TLS Proxy pane.

Step 1 Complete the first two steps of the Add TLS Proxy Instance Wizard. See Adding a TLS Proxy Instance, page 25-135 and Add TLS Proxy Instance Wizard – Client Configuration, page 25-137.

The Add TLS Proxy Instance Wizard – Client Configuration dialog box opens.

Step 2 To specify a client proxy certificate to use for the TLS Proxy, perform the following. Select this option when the client proxy certificate is being used between two servers; for example, when configuring the TLS Proxy for Presence Federation, which uses the Cisco Unified Presence Server (CUPS), both the TLS client and TLS server are both servers.

a. Check the Specify the proxy certificate for the TLS Client... check box.

b. Select a certificate from the drop-down list.

Or

To create a new client proxy certificate, click **Manage**. The Manage Identify Certificates dialog box opens. See Configuring Identity Certificates Authentication, page 34-8.

Note When you are configuring the TLS Proxy for the Phone Proxy and it is using the mixed security mode for the CUCM cluster, you must configure the LDC Issuer. The LDC Issuer lists the local certificate authority to issue client or server dynamic certificates.

Step 3 To specify an LDC Issuer to use for the TLS Proxy, perform the following. When you select and configure the LDC Issuer option, the security appliance acts as the certificate authority and issues certificates to TLS clients.

a. Click the Specify the internal Certificate Authority to sign the local dynamic certificate for phones... check box.

b. Click the Certificates radio button and select a self-signed certificate from the drop-down list or click **Manage** to create a new LDC Issuer. The Manage Identify Certificates dialog box opens. See Configuring Identity Certificates Authentication, page 34-8.

Or

Click the Certificate Authority radio button to specify a Certificate Authority (CA) server. When you specify a CA server, it needs to be created and enabled in the security appliance. To create and enable the CA server, click **Manage**. The Edit CA Server Settings dialog box opens. See Authenticating Using the Local CA, page 34-16.

Note To make configuration changes after the local certificate authority has been configured for the first time, disable the local certificate authority.

c. In the Key-Pair Name field, select a key pair from the drop-list. The list contains the already defined RSA key pair used by client dynamic certificates. To see the key pair details, including generation time, usage, modulus size, and key data, click **Show**.

Or

To create a new key pair, click **New**. The Add Key Pair dialog box opens. See Adding or Importing an Identity Certificate, page 34-9 for details about the Key Pair fields.

Step 4 In the Security Algorithms area, specify the available and active algorithms to be announced or matched during the TLS handshake.

- Available Algorithms—Lists the available algorithms to be announced or matched during the TLS handshake: des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, and null-sha1.

  Add—Adds the selected algorithm to the active list.

  Remove—Removes the selected algorithm from the active list.

- Active Algorithms—Lists the active algorithms to be announced or matched during the TLS handshake: des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, and null-sha1. For client proxy (acting as a TLS client to the server), the user-defined algorithms replace the original ones from the hello message for asymmetric encryption method between the two TLS legs. For example, the leg between the proxy and Call Manager may be NULL cipher to offload the Call Manager.

  Move Up—Moves an algorithm up in the list.

  Move Down—Moves an algorithm down in the list.

Step 5 Click **Next**.

The Add TLS Proxy Instance Wizard – Other Steps dialog box opens. The Other Steps dialog box provides instructions on the steps to complete outside the ASDM to make the TLS Proxy fully functional (see Add TLS Proxy Instance Wizard – Other Steps, page 25-139).

Add TLS Proxy Instance Wizard – Other Steps

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

The last dialog box of the Add TLS Proxy Instance Wizard specifies the additional steps required to make TLS Proxy fully functional. In particular, you need to perform the following tasks to complete the TLS Proxy configuration:

- Export the local CA certificate or LDC Issuer and install them on the original TLS server.

  To export the LDC Issuer, go to Configuration > Firewall > Advanced > Certificate Management > Identity Certificates > Export. See Exporting an Identity Certificate, page 34-11.
- For the TLS Proxy, enable Skinny and SIP inspection between the TLS server and TLS clients. See SIP Inspection, page 25-23 and Skinny (SCCP) Inspection, page 25-25. When you are configuring the TLS Proxy for Presence Federation (which uses CUP), you only enable SIP inspection because the feature supports only the SIP protocol.
- For the TLS Proxy for CUMA, enable MMP inspection. See MMP Inspection, page 25-20.
- When using the internal Certificate Authority of the security appliance to sign the LDC Issuer for TLS clients, perform the following:
  - Use the Cisco CTL Client to add the server proxy certificate to the CTL file and install the CTL file on the security appliance.

    For information on the Cisco CTL Client, see “Configuring the Cisco CTL Client” in *Cisco Unified CallManager Security Guide*.

    http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/security/5_0_4/secuauth.html

    To install the CTL file on the security appliance, go to Configuration > Firewall > Advanced > Encrypted Traffic Inspection > CTL Provider > Add. The Add CTL Provider dialog box opens. For information on using this dialog box to install the CTL file, see Add/Edit CTL Provider, page 25-153.
  - Create a CTL provider instance for connections from the CTL clients. See Add/Edit CTL Provider, page 25-153.

Edit TLS Proxy Instance – Server Configuration

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

The TLS Proxy enables inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and to support the Cisco Unified Communications features on the security appliance. For a detailed overview of the TLS Proxy used by these features, see TLS Proxy Wizard, page 25-133.

Use the Edit TLS Proxy – Server Configuration tab to edit the server proxy parameters for the original TLS Server—the Cisco Unified Call Manager (CUCM) server, the Cisco Unified Presence Server (CUPS), or the Cisco Unified Mobility Advantage (CUMA) server.

Step 1 Open the Configuration > Firewall > Advanced > Encrypted Traffic Inspection > TLS Proxy pane.

Step 2 To edit a TLS Proxy Instance, click **Edit**.

The Edit TLS Proxy Instance dialog box opens.

Step 3 If necessary, click the **Server Configuration** tab.

Step 4 Specify the server proxy certificate by doing one of the following:

- To add a new certificate, click **Manage**. The Manage Identify Certificates dialog box opens.

  When the Phone Proxy is operating in a mixed-mode CUCM cluster, you must import the CUCM certificate by clicking **Add** in the Manage Identify Certificates dialog box. See Adding or Importing an Identity Certificate, page 34-9.

- To select an existing certificate, select one from the drop-down list.

  When you are configuring the TLS Proxy for the Phone Proxy, select the certificate that has a filename beginning with **_internal_PP_**. When you create the CTL file for the Phone Proxy, the security appliance, creates an internal trustpoint used by the Phone Proxy to sign the TFTP files. The trustpoint is named **_internal_PP_***ctl-instance_filename*.

The server proxy certificate is used to specify the trustpoint to present during the TLS handshake. The trustpoint can be self-signed or enrolled locally with the certificate service on the proxy. For example, for the Phone Proxy, the server proxy certificate is used by the Phone Proxy during the handshake with the IP phones.

Step 5 To install the TLS server certificate in the security appliance trust store, so that the security appliance can authenticate the TLS server during TLS handshake between the proxy and the TLS server, click **Install TLS Server's Certificate**.

The Manage CA Certificates dialog box opens. See Guidelines and Limitations, page 34-2. Click **Add** to open the Install Certificate dialog box. See Adding or Importing an Identity Certificate, page 34-9.

When you are configuring the TLS Proxy for the Phone Proxy, click **Install TLS Server's Certificate** and install the Cisco Unified Call Manager (CUCM) certificate so that the proxy can authenticate the IP phones on behalf of the CUCM server.

Step 6 To require the security appliance to present a certificate and authenticate the TLS client during TLS handshake, check the Enable client authentication during TLS Proxy handshake check box.

When adding a TLS Proxy Instance for Mobile Advantage (the CUMC client and CUMA server), disable the check box when the client is incapable of sending a client certificate.

See TLS Proxy Wizard, page 25-133 to determine which TLS clients used by the Cisco Unified Communication features are capable of client authentication.

Step 7 Click **Apply** to save the changes.

Edit TLS Proxy Instance – Client Configuration

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

The TLS Proxy enables inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco Call Manager and to support the Cisco Unified Communications features on the security appliance. For a detailed overview of the TLS Proxy used by these features, see TLS Proxy Wizard, page 25-133.

The fields in the Edit TLS Proxy dialog box are identical to the fields displayed when you add a TLS Proxy instance. Use the Edit TLS Proxy – Client Configuration tab to edit the client proxy parameters for the original TLS Client, such as IP phones, CUMA clients, the Cisco Unified Presence Server (CUPS), or the Microsoft OCS server.

Step 1 Open the Configuration > Firewall > Advanced > Encrypted Traffic Inspection > TLS Proxy pane.

Step 2 To edit a TLS Proxy Instance, click **Edit**.

The Edit TLS Proxy Instance dialog box opens.

Step 3 If necessary, click the **Client Configuration** tab.

Step 4 To specify a client proxy certificate to use for the TLS Proxy, perform the following. Select this option when the client proxy certificate is being used between two servers; for example, when configuring the TLS Proxy for Presence Federation, which uses the Cisco Unified Presence Server (CUPS), both the TLS client and TLS server are both servers.

a. Check the Specify the proxy certificate for the TLS Client... check box.

b. Select a certificate from the drop-down list.

Or

To create a new client proxy certificate, click **Manage**. The Manage Identify Certificates dialog box opens. See Authenticating Using the Local CA, page 34-16.

Note When you are configuring the TLS Proxy for the Phone Proxy and it is using the mixed security mode for the CUCM cluster, you must configure the LDC Issuer. The LDC Issuer lists the local certificate authority to issue client or server dynamic certificates.

Step 5 To specify an LDC Issuer to use for the TLS Proxy, perform the following. When you select and configure the LDC Issuer option, the security appliance acts as the certificate authority and issues certificates to TLS clients.

a. Click the Specify the internal Certificate Authority to sign the local dynamic certificate for phones... check box.

b. Click the Certificates radio button and select a self-signed certificate from the drop-down list or click **Manage** to create a new LDC Issuer. The Manage Identify Certificates dialog box opens. See Authenticating Using the Local CA, page 34-16.

Or

Click the Certificate Authority radio button to specify a Certificate Authority (CA) server. When you specify a CA server, it needs to be created and enabled in the security appliance. To create and enable the CA server, click **Manage**. The Edit CA Server Settings dialog box opens. See Authenticating Using the Local CA, page 34-16.

Note To make configuration changes after the local certificate authority has been configured for the first time, disable the local certificate authority.

c. In the Key-Pair Name field, select a key pair from the drop-list. The list contains the already defined RSA key pair used by client dynamic certificates. To see the key pair details, including generation time, usage, modulus size, and key data, click **Show**.

Or

To create a new key pair, click **New**. The Add Key Pair dialog box opens. See Adding or Importing an Identity Certificate, page 34-9 for details about the Key Pair fields.

Step 6 In the Security Algorithms area, specify the available and active algorithms to be announced or matched during the TLS handshake.

- Available Algorithms—Lists the available algorithms to be announced or matched during the TLS handshake: des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, and null-sha1.

  Add—Adds the selected algorithm to the active list.

  Remove—Removes the selected algorithm from the active list.

- Active Algorithms—Lists the active algorithms to be announced or matched during the TLS handshake: des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, and null-sha1. For client proxy (acting as a TLS client to the server), the user-defined algorithms replace the original ones from the hello message for asymmetric encryption method between the two TLS legs. For example, the leg between the proxy and Call Manager may be NULL cipher to offload the Call Manager.

  Move Up—Moves an algorithm up in the list.

  Move Down—Moves an algorithm down in the list.

Step 7 Click **Apply** to save the changes.

Phone Proxy

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

For information on how to configure the Phone Proxy, see the following sections:

- Configuring the Phone Proxy, page 25-142
- Media Termination Instance Requirements, page 25-145
- Configuring the Media Termination Address, page 25-146
- Add/Edit TFTP Server, page 25-147

Use the Phone Proxy to configure a Phone Proxy between a Call Manager and IP phones. If the Phone Proxy is configured, the security appliance encrypts signaling connections from IP phones in the untrusted networks and sends them in the clear to the CUCM on a trusted network.

Configuring the Phone Proxy

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Configuring the Phone Proxy requires the following steps:

Step 1: Create the CTL file. See Creating a CTL File, page 25-148.

Step 2: Create the TLS Proxy instance to handle the encrypted signaling. See Adding a TLS Proxy Instance, page 25-135.

Step 3: Create the Phone Proxy instance. See Creating a Phone Proxy Instance, page 25-143.

Step 4: Configure the media termination address for the Phone Proxy. See Configuring the Media Termination Address, page 25-146.

Note Before you enable SIP and Skinny inspection for the Phone Proxy (which is done by applying the Phone Proxy to a service policy rule), the Phone Proxy must have an MTA instance, TLS Proxy, and CTL file assigned to it before the Phone Proxy can be applied to a service policy. Additionally, once a Phone Proxy is applied to a service policy rule, the Phone Proxy cannot be changed or removed.

Step 5: Enable the Phone Proxy with SIP and Skinny inspection. See SIP Inspection, page 25-23 and Skinny (SCCP) Inspection, page 25-25.

Creating a Phone Proxy Instance

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Use the Configure Phone Proxy pane to add a Phone Proxy. For a detailed overview of the Phone Proxy used by the security appliance, see Phone Proxy, page 25-142.

This pane is available from the Configuration > Firewall > Advanced > Encrypted Traffic Inspection > Phone Proxy pane.

Step 1 Open the Configuration > Firewall > Advanced > Encrypted Traffic Inspection > Phone Proxy pane.

Step 2 Check the Enable Phone Proxy check box to enable the feature.

Step 3 Check the Add MTA instance to Phone Proxy check box to add the media termination address to the Phone Proxy instance. You must have a media termination address instance configured. The configured address is added to the Phone Proxy instance. See Configuring the Media Termination Address, page 25-146.

Note To configure the media termination address, click the Configure MTA button. The Media Termination Address dialog box appears. Once you click the Add MTA instance to Phone Proxy check box, the media termination address instance cannot be modified and the button changes to View MTA Configuration. To change the media termination address, uncheck the Add MTA instance to Phone Proxy check box.

Step 4 Specify the TLS Proxy by doing one of the following:

- To add a new TLS Proxy Instance, click **Manage**. The Configure TLS Proxy dialog box opens. See Configure TLS Proxy Pane, page 25-134.
- To select an existing TLS Proxy, select one from the drop-down list.

Step 5 In the TFTP Server Settings list, do one of the following:

- To add a new TFTP server for the Phone Proxy, click **Add**. The Add TFTP Server dialog box opens. See Add/Edit TFTP Server, page 25-147.
- To select an existing TFTP server, select one from the drop-down list.

Note The TFTP server must reside on the same interface as the Cisco Unified Call Manager. Additionally, If NAT is configured for the TFTP server, the NAT configuration must be configured prior to configuring the specifying the TFTP server while creating the Phone Proxy instance.

Step 6 Specify the CTL File to use for the Phone Proxy by doing one of the following:

- To use an existing CTL File, check the Use the Certificate Trust List File generated by the CTL instance check box.
- To create a new CTL file for the Phone Proxy, click the link Generate Certificate Trust List File. The Create a Certificate Trust List (CTL) File pane opens. See Creating a CTL File, page 25-148.

Step 7 To specify the security mode of the CUCM cluster, click one of the following options in the CUCM Cluster Mode field:

- Non-secure—Specifies the cluster mode to be in nonsecure mode when configuring the Phone Proxy feature.
- Mixed—Specifies the cluster mode to be in mixed mode when configuring the Phone Proxy feature.

Step 8 To configure the idle timeout after which the secure-phone entry is removed from the Phone Proxy database (the default is 5 minutes), enter a value in the format *hh*:*mm*:*ss*.

Since secure phones always request a CTL file upon bootup, the Phone Proxy creates a database that marks the phone as secure. The entries in the secure phone database are removed after a specified configured timeout. The entry timestamp is updated for each registration refresh the Phone Proxy receives for SIP phones and KeepAlives for SCCP phones.

Specify a value that is greater than the maximum timeout value for SCCP KeepAlives and SIP Register refresh. For example, if the SCCP KeepAlives are configured for 1 minute intervals and the SIP Register Refresh is configured for 3 minutes, configure this timeout value greater than 3 minutes.

Step 9 To preserve Call Manager configuration on the IP phones, check the Preserve the Call Manager's configuration on the phone... When this option is uncheck, the following service settings are disabled on the IP phones:

- PC Port
- Gratuitous ARP
- Voice VLAN access
- Web Access
- Span to PC Port

Step 10 To configure an HTTP proxy for the Phone Proxy feature that is written into the IP phone's configuration file under the <proxyServerURL> tag, do the following:

a. Check the Configure a http-proxy which would be written into the phone's config file... check box.

b. In the IP Address field, type the IP address of the HTTP proxy and the listening port of the HTTP proxy.

The IP address you enter should be the global IP address based on where the IP phone and HTTP proxy server is located. You can enter a hostname in the IP Address field when that hostname can be resolved to an IP address by the security appliance (for example, DNS lookup is configured) because the security appliance will resolve the hostname to an IP address. If a port is not specified, the default will be 8080.

c. In the Interface field, select the interface on which the HTTP proxy resides on the security appliance.

Setting the proxy server configuration option for the Phone Proxy allows for an HTTP proxy on the DMZ or external network in which all the IP phone URLs are directed to the proxy server for services on the phones. This setting accommodates nonsecure HTTP traffic, which is not allowed back into the corporate network.

Step 11 To force Cisco IP Communicator (CIPC) softphones to operate in authenticated mode when CIPC softphones are deployed in a voice and data VLAN scenario, check the Enable CIPC security mode authentication check box.

Because CIPC requires an LSC to perform the TLS handshake, CIPC needs to register with the CUCM in nonsecure mode using cleartext signaling. To allow the CIPC to register, create an ACL that allows the CIPC to connect to the CUCM on the nonsecure SIP/SCCP signalling ports (5060/2000).

CIPC uses a different cipher when doing the TLS handshake and requires the null-sha1 cipher and SSL encryption be configured. To add the null-shal cipher, go to Configuration > Device Management > Advanced > SSL Settings > Encryption section. Select the null-shal SSL encryption type and add it to the Available Algorithms.

Current versions of Cisco IP Communicator (CIPC) support authenticated mode and perform TLS signaling but not voice encryption.

Step 12 Click **Apply** to save the Phone Proxy configuration settings.

Note After creating the Phone Proxy instance, you enable it with SIP and Skinny inspection. See SIP Inspection, page 25-23 and Skinny (SCCP) Inspection, page 25-25.

However, before you enable SIP and Skinny inspection for the Phone Proxy (which is done by applying the Phone Proxy to a service policy rule), the Phone Proxy must have an MTA instance, TLS Proxy, and CTL file assigned to it before the Phone Proxy can be applied to a service policy. Additionally, once a Phone Proxy is applied to a service policy rule, the Phone Proxy cannot be changed or removed.

Media Termination Instance Requirements

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

The security appliance must have a media termination instance that meets the following criteria:

- You must configure one media termination instance for each phone proxy on the security appliance. Multiple media termination instances on the security appliance are not supported.
- For the media termination instance, you can configure a global media-termination address for all interfaces or configure a media-termination address for different interfaces. However, you cannot use a global media-termination address and media-termination addresses configured for each interface at the same time.
- If you configure a media termination address for multiple interfaces, you must configure an address on each interface that the security appliance uses when communicating with IP phones.

  For example, if you had three interfaces on the security appliance (one internal interface and two external interfaces) and only one of the external interfaces were used to communicate with IP phones, you would configure two media termination addresses: one on the internal interface and one on the external interface that communicated with the IP phones.

- Only one media-termination address can be configured per interface.
- The IP addresses are publicly routable addresses that are unused IP addresses within the address range on that interface.
- The IP address on an interface cannot be the same address as that interface on the security appliance.
- The IP addresses cannot overlap with existing static NAT pools or NAT rules.
- The IP addresses cannot be the same as the Cisco UCM or TFTP server IP address.
- For IP phones behind a router or gateway, you must also meet this prerequisite. On the router or gateway, add routes to the media termination address on the security appliance interface that the IP phones communicate with so that the phone can reach the media termination address.

Configuring the Media Termination Address

The media termination address you configure must meet the requirements as described in Media Termination Instance Requirements, page 25-145.

Note In versions before 8.2(1), you configured one media-termination address (MTA) on the outside interface of the adaptive security appliance where the remote Cisco IP phones were located. In Version 8.2(1) and later, you can configure a global media-termination address for all interfaces or configure a media-termination address for different interfaces.

As a result of this enhancement, the old configuration has been deprecated. You can continue to use the old configuration if desired. However, if you need to change the configuration at all, only the new configuration method is accepted; you cannot later restore the old configuration. If you need to maintain downgrade compatibility, you should keep the old configuration as is.

Step 1 Open the Configuration > Firewall > Advanced > Encrypted Traffic Inspection > Media Termination Address pane.

Step 2 In the Media Termination Address Settings section, specify whether to configure a media-termination address (MTA) per interface or to configure a global MTA. You can configure a global media-termination address for all interfaces or configure a media-termination address for different interfaces.

- To configure an MTA per interface, click the Configure MTA per Interface radio button and click the **Add** button. In the dialog box that appears, specify the interface name and enter an IP address or hostname.

  If you configure a media termination address for multiple interfaces, you must configure an address on each interface that the security appliance uses when communicating with IP phones. The IP addresses are publicly routable addresses that are unused IP addresses within the address range on that interface.

  See Media Termination Instance Requirements, page 25-145 for the complete list of requirements that you must follow when creating the media termination instance and configuring the media termination addresses.

- To configure a global MTA, click the Configure Global MTA radio button and enter the IP address in the text box. See Media Termination Instance Requirements, page 25-145 for the complete list of requirements that you must follow when configuring a global media termination address.

Step 3 Specify the minimum and maximum values for the RTP port range for the media termination instance. The minimum port can be a value from 1024 to 16384 and the maximum port can be a value from 32767 to 65535.

Step 4 To apply the media termination address you are configuring to the existing Phone Proxy instance, click the Apply MTA to Phone Proxy Service Policy check box.

Step 5 Click **Apply** to save the media termination address configuration settings.

Add/Edit TFTP Server

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Use the Add/Edit TFTP Server dialog box to specify the IP address of the TFTP server and the interface on which the TFTP server resides.

The Phone Proxy must have at least one CUCM TFTP server configured. Up to five TFTP servers can be configured for the Phone Proxy.

The TFTP server is assumed to be behind the firewall on the trusted network; therefore, the Phone Proxy intercepts the requests between the IP phones and TFTP server.

Note If NAT is configured for the TFTP server, the NAT configuration must be configured prior to specifying the TFTP server while creating the Phone Proxy instance.

Fields

TFTP Server IP Address—Specifies the address of the TFTP server. Create the TFTP server using the actual internal IP address.

Port—(Optional) Specifies the port the TFTP server is listening in on for the TFTP requests. This should be configured if it is not the default TFTP port 69.

Interface—Specifies the interface on which the TFTP server resides. The TFTP server must reside on the same interface as the Cisco Unified Call Manager (CUCM).

Modes

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

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

CTL File

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

For information on how to configure CTL files, see the following sections:

- Creating a CTL File, page 25-148
- Add/Edit Record Entry, page 25-149
- CTL Provider, page 25-152

Create a Certificate Trust List (CTL) file that is required by the Phone Proxy. Specify the certificates needed by creating a new CTL file or by specifying the path of an exiting CTL file to parse from Flash memory.

Create trustpoints and generate certificates for each entity in the network (CUCM, CUCM and TFTP, TFTP server, CAPF) that the IP phones must trust. The certificates are used in creating the CTL file. You need to create trustpoints for each CUCM (primary and secondary if a secondary CUCM is used) and TFTP server in the network. The trustpoints need to be in the CTL file for the phones to trust the CUCM.

Create the CTL File that will be presented to the IP phones during the TFTP. The address must be the translated or global address of the TFTP server or CUCM if NAT is configured.

When the file is created, it creates an internal trustpoint used by the Phone Proxy to sign the TFTP files. The trustpoint is named **_internal_PP_***ctl-instance_filename*.

Creating a CTL File

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Note When a CTL file instance is assigned to the Phone Proxy, you cannot modify it in the CTL File pane and the pane is disabled. To modify a CTL File that is assigned to the Phone Proxy, go to the Phone Proxy pane (Configuration > Firewall > Advanced > Encrypted Traffic Inspection > Phone Proxy), and deselect the Use the Certificate Trust List File generated by the CTL instance check box.

Use the Create a Certificate Trust List (CTL) File pane to create a CTL file for the Phone Proxy. This pane creates the CTL file that is presented to the IP phones during the TFTP handshake with the security appliance. For a detailed overview of the CTL file used by the Phone Proxy, see CTL File, page 25-147.

The Create a Certificate Trust List (CTL) File pane is used to configure the attributes for generating the CTL file. The name of the CTL file instance is generated by the ASDM. When the user tries to edit the CTL file instance configuration, the ASDM automatically generates the **shutdown** CLI command first and the **no shutdown** CLI command as the last command.

This pane is available from the Configuration > Firewall > Advanced > Encrypted Traffic Inspection > CTL File pane.

Step 1 Open the Configuration > Firewall > Advanced > Encrypted Traffic Inspection > CTL File pane.

Step 2 Check the Enable Certificate Trust List File check box to enable the feature.

Step 3 To specify the CTL file to use for the Phone Proxy, perform one of the following:

- If there is an existing CTL file available, download the CTL file to Flash memory by using the File Management Tool in the ASDM Tools menu. Select the Use certificates present in the CTL stored in flash radio button and specify the CTL file name and path in the text box.

Use an existing CTL file to install the trustpoints for each entity in the network (CUCM, CUCM and TFTP, TFTP server, CAPF) that the IP phones must trust. If you have an existing CTL file that contains the correct IP addresses of the entities (namely, the IP address that the IP phones use for the CUCM or TFTP servers), you can be use it to create a new CTL file. Store a copy of the existing CTL file to Flash memory and rename it something other than `CTLFile.tlv`

- If there is no existing CTL file available, select Create new CTL file radio button.

  Add Record entries for each entity in the network such as CUCM, TFTP, and CUCM-TFTP option by clicking **Add**. The Add Record Entry dialog box opens. See Add/Edit Record Entry, page 25-149.

Step 4 Specify the number SAST certificate tokens required. The default is 2. maximum allowed is 5.

Because the Phone Proxy generates the CTL file, it needs to create the System Administrator Security Token (SAST) key to sign the CTL file itself. This key can be generated on the security appliance. A SAST is created as a self-signed certificate. Typically, a CTL file contains more than one SAST. In case a SAST is not recoverable, the other one can be used to sign the file later.

Step 5 Click **Apply** to save the CTL file configuration settings.

Add/Edit Record Entry

Note This feature is not supported for the Adaptive Security Appliance version 8.1.2.

Use the Add/Edit Record Entry dialog box to specify the trustpoints to be used for the creation of the CTL file.

Add additional record-entry configurations for each entity that is required in the CTL file.

Fields

Type—Specifies the type of trustpoint to create:

- cucm: Specifies the role of this trustpoint to be CCM. Multiple CCM trustpoints can be configured.
- cucm-tftp: Specifies the role of this trustpoint to be CCM+TFTP. Multiple CCM+TFTP trustpoints can be configured.
- tftp: Specifies the role of this trustpoint to be TFTP. Multiple TFTP trustpoints can be configured.
- capf: Specifies the role of this trustpoint to be CAPF. Only one CAPF trustpoint can be configured.

Address—Specifies the IP address of the trustpoint. The IP address you specify must be the global address of the TFTP server or CUCM if NAT is configured. The global IP address is the IP address as seen by the IP phones because it will be the IP address used for the CTL record for the trustpoint.

Certificate—Specifies the Identity Certificate for the record entry in the CTL file. You can create a new Identity Certificate by clicking **Manage**. The Manage Identify Certificates dialog box opens. See Authenticating Using the Local CA, page 34-16.

You can add an Identity Certificate by generating a self-signed certificate, obtaining the certificate through SCEP enrollment, or by importing a certificate in PKCS-12 format. Choose the best option based on the requirements for configuring the CTL file.

Domain Name—(Optional) Specifies the domain name of the trustpoint used to create the DNS field for the trustpoint. This is appended to the Common Name field of the Subject DN to create the DNS Name. The domain name should be configured when the FQDN is not configured for the trustpoint. Only one domain-name can be specified.

Note If you are using domain names for your CUCM and TFTP server, you must configure DNS lookup on the security appliance. Add an entry for each of the outside interfaces on the security appliance into your DNS server, if such entries are not already present. Each security appliance outside IP address should have a DNS entry associated with it for lookups. These DNS entries must also be enabled for Reverse Lookup. Additionally, define your DNS server IP address on the security appliance; for example: `dns name-server 10.2.3.4` (IP address of your DNS server).

Modes

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

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

TLS Proxy

that feature is supported only for ASA versions 8.0.x prior to 8.0.4 and for version 8.1.

Note This feature is not supported for the Adaptive Security Appliance versions prior to 8.0.4 and for version 8.1.2.

Use the TLS Proxy option to enable inspection of SSL encrypted VoIP signaling, namely Skinny and SIP, interacting with Cisco CallManager.

The TLS Proxy pane lets you define and configure Transaction Layer Security Proxy to enable inspection of encrypted traffic.

Fields

- TLS Proxy Name—Lists the TLS Proxy name.
- Server—Lists the trustpoint, which is either self-signed or enrolled with a certificate server.
- Local Dynamic Certificate Issuer—Lists the local certificate authority to issue client or server dynamic certificates.
- Local Dynamic Certificate Key Pair—Lists the RSA key pair used by client or server dynamic certificates.
- Add—Adds a TLS Proxy.
- Edit—Edits a TLS Proxy.
- Delete—Deletes a TLS Proxy.
- Maximum Sessions—Lets you specify the maximum number of TLS Proxy sessions to support.
  - Specify the maximum number of TLS Proxy sessions that the ASA needs to support. By default, ASA supports 300 sessions.—Enables maximum number of sessions option.
  - Maximum number of sessions:—The minimum is 1. The maximum is dependent on the platform. The default is 300.

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 TLS Proxy

Note This feature is not supported for the Adaptive Security Appliance versions prior to 8.0.4 and for version 8.1.2.

The Add/Edit TLS Proxy dialog box lets you define the parameters for the TLS Proxy.

Fields

- TLS Proxy Name—Specifies the TLS Proxy name.
- Server Configuration—Specifies the proxy certificate name.
  - Server—Specifies the trustpoint to be presented during the TLS handshake. The trustpoint could be self-signed or enrolled locally with the certificate service on the proxy.
- Client Configuration—Specifies the local dynamic certificate issuer and key pair.
  - Local Dynamic Certificate Issuer—Lists the local certificate authority to issue client or server dynamic certificates.

    Certificate Authority Server—Specifies the certificate authority server.

    Certificate—Specifies a certificate.

    Manage—Configures the local certificate authority. To make configuration changes after it has been configured for the first time, disable the local certificate authority.
  - Local Dynamic Certificate Key Pair—Lists the RSA key pair used by client dynamic certificates.

    Key-Pair Name—Specifies a defined key pair.

    Show—Shows the key pair details, including generation time, usage, modulus size, and key data.

    New—Lets you define a new key pair.
- More Options—Specifies the available and active algorithms to be announced or matched during the TLS handshake.
  - Available Algorithms—Lists the available algorithms to be announced or matched during the TLS handshake: des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, and null-sha1.

    Add—Adds the selected algorithm to the active list.

    Remove—Removes the selected algorithm from the active list.

  - Active Algorithms—Lists the active algorithms to be announced or matched during the TLS handshake: des-sha1, 3des-sha1, aes128-sha1, aes256-sha1, and null-sha1. For client proxy (acting as a TLS client to the server), the user-defined algorithms replace the original ones from the hello message for asymmetric encryption method between the two TLS legs. For example, the leg between the proxy and CallManager may be NULL cipher to offload the CallManager.

    Move Up—Moves an algorithm up in the list.

    Move Down—Moves an algorithm down in the list.

Modes

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

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

CTL Provider

Use the CTL Provider option to configure Certificate Trust List provider service.

The CTL Provider pane lets you define and configure Certificate Trust List provider service to enable inspection of encrypted traffic.

Fields

- CTL Provider Name—Lists the CTL Provider name.
- Client Details—Lists the name and IP address of the client.
  - Interface Name—Lists the defined interface name.
  - IP Address—Lists the defined interface IP address.
- Certificate Name—Lists the certificate to be exported.
- Add—Adds a CTL Provider.
- Edit—Edits a CTL Provider.
- Delete—Deletes a CTL Provider.

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 CTL Provider

The Add/Edit CTL Provider dialog box lets you define the parameters for the CTL Provider.

Fields

- CTL Provider Name—Specifies the CTL Provider name.
- Certificate to be Exported—Specifies the certificate to be exported to the client.
  - Certificate Name—Specifies the name of the certificate to be exported to the client.
  - Manage—Manages identity certificates. See Authenticating Using the Local CA, page 34-16
- Client Details—Specifies the clients allowed to connect.
  - Client to be Added—Specifies the client interface and IP address to add to the client list.

    Interface—Specifies client interface.

    IP Address—Specifies the client IP address.

    Add—Adds the new client to the client list.

    Delete—Deletes the selected client from the client list.
- More Options—Specifies the available and active algorithms to be announced or matched during the TLS handshake.
  - Parse the CTL file provided by the CTL Client and install trustpoints—Trustpoints installed by this option have names prefixed with "_internal_CTL_." If disabled, each Call Manager server and CAPF certificate must be manually imported and installed.
  - Port Number—Specifies the port to which the CTL provider listens. The port must be the same as the one listened to by the CallManager servers in the cluster (as configured under Enterprise Parameters on the CallManager administration page). The default is 2444.
  - Authentication—Specifies the username and password that the client authenticates with the provider.

    Username—Client username.

    Password—Client password.

    Confirm Password—Client password.

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

Configuring QoS

Have you ever participated in a long-distance phone call that involved a satellite connection? The conversation might be interrupted with brief, but perceptible, gaps at odd intervals. Those gaps are the time, called the latency, between the arrival of packets being transmitted over the network. Some network traffic, such as voice and video, cannot tolerate long latency times. Quality of Service (QoS) is a feature that lets you give priority to critical traffic, prevent bandwidth hogging, and manage network bottlenecks to prevent packet drops.

This chapter describes how to apply QoS policies, and includes the following sections:

QoS Overview

You should consider that in an ever-changing network environment, QoS is not a one-time deployment, but an ongoing, essential part of network design.

Note QoS is only available in single context mode.

This section describes the QoS features supported by the security appliance, and includes the following topics:

Supported QoS Features

The security appliance supports the following QoS features:

- Policing—To prevent individual flows from hogging the network bandwidth, you can limit the maximum bandwidth used per flow. See the "Policing Overview" section on page 26-3 for more information.
- Priority queuing—For critical traffic that cannot tolerate latency, such as Voice over IP (VoIP), you can identify traffic for Low Latency Queuing (LLQ) so that it is always transmitted ahead of other traffic. See the "Priority Queueing Overview" section on page 26-3 for more information.
- Traffic shaping—If you have a device that transmits packets at a high speed, such as a security appliance with Fast Ethernet, and it is connected to a low speed device such as a cable modem, then the cable modem is a bottleneck at which packets are frequently dropped. To manage networks with differing line speeds, you can configure the security appliance to transmit packets at a fixed slower rate. See the "Traffic Shaping Overview" section on page 26-4 for more information.

What is a Token Bucket?

A token bucket is used to manage a device that regulates the data in a flow. For example, the regulator might be a traffic policer or a traffic shaper. A token bucket itself has no discard or priority policy. Rather, a token bucket discards tokens and leaves to the flow the problem of managing its transmission queue if the flow overdrives the regulator.

A token bucket is a formal definition of a rate of transfer. It has three components: a burst size, an average rate, and a time interval. Although the average rate is generally represented as bits per second, any two values may be derived from the third by the relation shown as follows:

average rate = burst size / time interval

These terms are defined as follows:

- Average rate—Also called the committed information rate (CIR), it specifies how much data can be sent or forwarded per unit time on average.
- Burst size—Also called the Committed Burst (Bc) size, it specifies in bits or bytes per burst how much traffic can be sent within a given unit of time to not create scheduling concerns. (For traffic shaping, it specifies bits per burst; for policing, it specifies bytes per burst.)
- Time interval—Also called the measurement interval, it specifies the time quantum in seconds per burst.

In the token bucket metaphor, tokens are put into the bucket at a certain rate. The bucket itself has a specified capacity. If the bucket fills to capacity, newly arriving tokens are discarded. Each token is permission for the source to send a certain number of bits into the network. To send a packet, the regulator must remove from the bucket a number of tokens equal in representation to the packet size.

If not enough tokens are in the bucket to send a packet, the packet either waits until the bucket has enough tokens (in the case of traffic shaping) or the packet is discarded or marked down (in the case of policing). If the bucket is already full of tokens, incoming tokens overflow and are not available to future packets. Thus, at any time, the largest burst a source can send into the network is roughly proportional to the size of the bucket.

Note that the token bucket mechanism used for traffic shaping has both a token bucket and a data buffer, or queue; if it did not have a data buffer, it would be a policer. For traffic shaping, packets that arrive that cannot be sent immediately are delayed in the data buffer.

For traffic shaping, a token bucket permits burstiness but bounds it. It guarantees that the burstiness is bounded so that the flow will never send faster than the token bucket capacity, divided by the time interval, plus the established rate at which tokens are placed in the token bucket. See the following formula:

(token bucket capacity in bits / time interval in seconds) + established rate in bps = maximum flow speed in bps

This method of bounding burstiness also guarantees that the long-term transmission rate will not exceed the established rate at which tokens are placed in the bucket.

Policing Overview

Policing is a way of ensuring that no traffic exceeds the maximum rate (in bits/second) that you configure, thus ensuring that no one traffic flow or class can take over the entire resource. When traffic exceeds the maximum rate, the security appliance drops the excess traffic. Policing also sets the largest single burst of traffic allowed.

Priority Queueing Overview

LLQ priority queueing lets you prioritize certain traffic flows (such as latency-sensitive traffic like voice and video) ahead of other traffic.

The security appliance supports two types of priority queueing:

- Standard priority queueing—Standard priority queueing uses an LLQ priority queue on an interface (see the "Creating the Standard Priority Queue for an Interface" section on page 26-5), while all other traffic goes into the "best effort" queue. Because queues are not of infinite size, they can fill and overflow. When a queue is full, any additional packets cannot get into the queue and are dropped. This is called *tail drop*. To avoid having the queue fill up, you can increase the queue buffer size. You can also fine-tune the maximum number of packets allowed into the transmit queue. These options let you control the latency and robustness of the priority queuing. Packets in the LLQ queue are always transmitted before packets in the best effort queue.
- Hierarchical priority queueing—Hierarchical priority queueing is used on interfaces on which you enable a traffic shaping queue. A subset of the shaped traffic can be prioritized. The standard priority queue is not used. See the following guidelines about hierarchical priority queueing:
  - Priority packets are always queued at the head of the shape queue so they are always transmitted ahead of other non-priority queued packets.
  - Priority packets are never dropped from the shape queue unless the sustained rate of priority traffic exceeds the shape rate.
  - For IPSec-encrypted packets, you can only match traffic based on the DSCP or precedence setting.
  - IPSec-over-TCP is not supported for priority traffic classification.

Traffic Shaping Overview

Traffic shaping is used to match device and link speeds, thereby controlling packet loss, variable delay, and link saturation, which can cause jitter and delay.

Note Traffic shaping is not supported on the ASA 5580.

- Traffic shaping must be applied to all outgoing traffic on a physical interface or in the case of the ASA 5505, on a VLAN. You cannot configure traffic shaping for specific types of traffic.
- Traffic shaping is implemented when packets are ready to be transmitted on an interface, so the rate calculation is performed based on the actual size of a packet to be transmitted, including all the possible overhead such as the IPSec header and L2 header.
- The shaped traffic includes both through-the-box and from-the-box traffic.
- The shape rate calculation is based on the standard token bucket algorithm. The token bucket size is twice the Burst Size value. See the "What is a Token Bucket?" section on page 26-2.
- When bursty traffic exceeds the specified shape rate, packets are queued and transmitted later. Following are some characteristics regarding the shape queue (for information about hierarchical priority queueing, see the "Priority Queueing Overview" section on page 26-3):
  - The queue size is calculated based on the shape rate. The queue can hold the equivalent of 200-milliseconds worth of shape rate traffic, assuming a 1500-byte packet. The minimum queue size is 64.
  - When the queue limit is reached, packets are tail-dropped.
  - Certain critical keep-alive packets such as OSPF Hello packets are never dropped.
  - The time interval is derived by *time_interval* = *burst_size* / *average_rate*. The larger the time interval is, the burstier the shaped traffic might be, and the longer the link might be idle. The effect can be best understood using the following exaggerated example:

    Average Rate = 1000000

    Burst Size = 1000000

    In the above example, the time interval is 1 second, which means, 1 Mbps of traffic can be bursted out within the first 10 milliseconds of the 1-second interval on a 100 Mbps FE link and leave the remaining 990 milliseconds idle without being able to send any packets until the next time interval. So if there is delay-sensitive traffic such as voice traffic, the Burst Size should be reduced compared to the average rate so the time interval is reduced.

How QoS Features Interact

You can configure each of the QoS features alone if desired for the security appliance. Often, though, you configure multiple QoS features on the security appliance so you can prioritize some traffic, for example, and prevent other traffic from causing bandwidth problems.

See the following supported feature combinations per interface:

- Standard priority queuing (for specific traffic) + Policing (for the rest of the traffic).

  You cannot configure priority queueing and policing for the same set of traffic.
- Traffic shaping (for all traffic on an interface) + Hierarchical priority queueing (for a subset of traffic).

You cannot configure traffic shaping and standard priority queueing for the same interface; only hierarchical priority queueing is allowed. For example, if you configure standard priority queueing for the global policy, and then configure traffic shaping for a specific interface, the feature you configured last is rejected because the global policy overlaps the interface policy.

Typically, if you enable traffic shaping, you do not also enable policing for the same traffic, although the security appliance does not restrict you from configuring this.

DSCP and DiffServ Preservation

- DSCP markings are preserved on all traffic passing through the security appliance.
- The security appliance does not locally mark/remark any classified traffic, but it honors the Expedited Forwarding (EF) DSCP bits of every packet to determine if it requires "priority" handling and will direct those packets to the LLQ.
- DiffServ marking is preserved on packets when they traverse the service provider backbone so that QoS can be applied in transit (QoS tunnel pre-classification).

Creating the Standard Priority Queue for an Interface

If you enable standard priority queueing for traffic on a physical interface, then you need to also create the priority queue on each interface. Each physical interface uses two queues: one for priority traffic, and the other for all other traffic. For the other traffic, you can optionally configure policing.

Note The standard priority queue is not required for hierarchical priority queueing with traffic shaping; see the "Priority Queueing Overview" section on page 26-3 for more information.

To create the priority queue, perform the following steps:

Step 1 Go to Configuration > Device Management > Advanced > Priority Queue, and click **Add**.

The Add Priority Queue dialog box displays.

Step 2 From the Interface drop-down list, choose the physical interface name on which you want to enable the priority queue, or for the ASA 5505, the VLAN interface name.

Step 3 To change the size of the priority queues, in the Queue Limit field, enter the number of average, 256-byte packets that the specified interface can transmit in a 500-ms interval.

A packet that stays more than 500 ms in a network node might trigger a timeout in the end-to-end application. Such a packet can be discarded in each network node.

Because queues are not of infinite size, they can fill and overflow. When a queue is full, any additional packets cannot get into the queue and are dropped (called *tail drop*). To avoid having the queue fill up, you can use this option to increase the queue buffer size.

The upper limit of the range of values for this option is determined dynamically at run time. The key determinants are the memory needed to support the queues and the memory available on the device.

The Queue Limit that you specify affects both the higher priority low-latency queue and the best effort queue.

Step 4 To specify the depth of the priority queues, in the Transmission Ring Limit field, enter the number of maximum 1550-byte packets that the specified interface can transmit in a 10-ms interval.

This setting guarantees that the hardware-based transmit ring imposes no more than 10-ms of extra latency for a high-priority packet.

This option sets the maximum number of low-latency or normal priority packets allowed into the Ethernet transmit driver before the driver pushes back to the queues on the interface to let them buffer packets until the congestion clears.

The upper limit of the range of values is determined dynamically at run time. The key determinants are the memory needed to support the queues and the memory available on the device.

The Transmission Ring Limit that you specify affects both the higher priority low-latency queue and the best-effort queue.

Creating a Policy for Standard Priority Queueing and/or Policing

You can configure standard priority queueing and policing rules for the same interface. See the "How QoS Features Interact" section on page 26-4 for information about valid QoS configurations.

To configure a QoS service policy, perform the following steps:

Step 1 To configure priority queueing, configure a service policy rule in the Configuration > Firewall > Service Policy Rules pane according to Chapter 23, "Configuring Service Policy Rules."

You can configure QoS as part of a new service policy rule, or you can edit an existing service policy.

For priority traffic, identify only latency-sensitive traffic. You can match traffic based on many characteristics, including access lists, tunnel groups, DSCP, precedence, and more. You cannot use the **class-default** class map for priority traffic. You cannot configure priority queueing for the global policy if you also enable traffic shaping on any interfaces.

Step 2 In the Rule Actions dialog box, click the **QoS** tab.

Step 3 Click **Enable priority for this flow**.

If this service policy rule is for an individual interface, ASDM automatically creates the priority queue for the interface (Configuration > Properties > Priority Queue; for more information, see the "Creating the Standard Priority Queue for an Interface" section on page 26-5). If this rule is for the global policy, then you need to manually add the priority queue to one or more interfaces *before* you configure the service policy rule.

Step 4 Click **Finish**. The service policy rule is added to the rule table.

Step 5 To configure policing, configure a service policy rule for the same interface in the Configuration > Firewall > Service Policy Rules pane according to Chapter 23, "Configuring Service Policy Rules."

For policing traffic, you can choose to police all traffic that you are not prioritizing, or you can limit the traffic to certain types.

Step 6 In the Rule Actions dialog box, click the **QoS** tab.

Step 7 Click **Enable policing**, then check the **Input policing** or **Output policing** (or both) check boxes to enable the specified type of traffic policing. For each type of traffic policing, configure the following fields:

- Committed Rate—The rate limit for this traffic flow; this is a value in the range 8000-2000000000, specifying the maximum speed (bits per second) allowed.

- Conform Action—The action to take when the rate is less than the conform-burst value. Values are transmit or drop.
- Exceed Action—Take this action when the rate is between the conform-rate value and the conform-burst value. Values are transmit or drop.
- Burst Rate—A value in the range 1000-512000000, specifying the maximum number of instantaneous bytes allowed in a sustained burst before throttling to the conforming rate value.

Step 8 Click **Finish**. The service policy rule is added to the rule table.

Step 9 Click **Apply** to send the configuration to the device.

Creating a Policy for Traffic Shaping and Hierarchical Priority Queueing

You can configure traffic shaping for all traffic on an interface, and optionally hierarchical priority queueing for a subset of latency-sensitive traffic. See the "How QoS Features Interact" section on page 26-4 for information about valid QoS configurations.

Note One side-effect of priority queueing is packet re-ordering. For IPSec packets, out-of-order packets that are not within the anti-replay window generate warning syslog messages. These warnings are false alarms in the case of priority queueing. You can configure the IPSec anti-replay window size to avoid possible false alarms. See the Configuration > VPN > IPSec > IPSec Rules > Enable Anti-replay window size option in the "Adding Crypto Maps" section on page 35-10.

Traffic shaping is not supported on the ASA 5580.

To configure a QoS service policy, perform the following steps:

Step 1 Configure a service policy on the Configuration > Firewall > Service Policy Rules pane according to Chapter 23, "Configuring Service Policy Rules."

You can configure QoS as part of a new service policy rule, or you can edit an existing service policy.

For traffic shaping, all traffic on an interface must be shaped. You can only use the **class-default** class map, which is automatically created by the security appliance, and which matches all traffic.

You cannot configure a separate traffic shaping rule on the same interface for which you configure a priority queueing rule (see the "Creating a Policy for Standard Priority Queueing and/or Policing" section on page 26-6); you can, however, configure priority queueing for a subset of shaped traffic under the traffic shaping rule. You also cannot configure traffic shaping for the global policy if you also enable priority queueing on any interfaces.

Step 2 In the Rule Actions dialog box, click the **QoS** tab.

Step 3 Click **Enable traffic shaping**, and configure the following fields:

- Average Rate—Sets the average rate of traffic in bits per second over a given fixed time period, between 64000 and 154400000. Specify a value that is a multiple of 8000.

- Burst Size—Sets the average burst size in bits that can be transmitted over a given fixed time period, between 2048 and 154400000. Specify a value that is a multiple of 128. If you do not specify the Burst Size, the default value is equivalent to 4-milliseconds of traffic at the specified Average Rate. For example, if the average rate is 1000000 bits per second, 4 ms worth = 1000000 * 4/1000 = 4000.

Step 4 (Optional) To configure priority queueing for a subset of shaped traffic:

a. Click **Enforce priority to selected shape traffic**.

b. Click **Configure** to identify the traffic that you want to prioritize.

   You are prompted to identify the traffic for which you want to apply priority queueing.

c. After you identify the traffic (see the "Adding a Service Policy Rule for Through Traffic" section on page 23-7), click **Next**.

d. Click **Enable priority for this flow**.

e. Click **Finish**.

   You return to the QoS tab.

Note For this type of priority queueing, you do *not* need to create a priority queue on an interface (**Configuration** > **Properties** > **Priority Queue**).

Step 5 Click **Finish**. The service policy rule is added to the rule table.

Step 6 Click **Apply** to send the configuration to the device.

CHAPTER 27

Configuring Filter Rules

This chapter describes ways to filter web traffic to reduce security risks or prevent inappropriate use. This chapter includes the following sections:

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.

This section includes the following topics:

Configuring URL Filtering

To enable filtering with an external filtering server, perform the following steps:

Step 1 Choose **Configuration** > **Firewall** > **URL Filter Servers** to specify an external filtering server. See URL Filtering Servers, page 27-2.

Step 2 (Optional) Buffer responses from the content server. See Advanced URL Filtering, page 27-4.

Step 3 (Optional) Cache content server addresses to improve performance. See Advanced URL Filtering, page 27-4.

Step 4 Choose **Configuration** > **Firewall** > **Filter Rules** to configure filter rules. See Filter Rules, page 27-5.

Step 5 Configure the external filtering server. For more information see the following websites:

- http://www.websense.com
- http://www.securecomputing.com

URL Filtering Servers

The URL Filtering Servers pane lets you specify the external filter server to use. You can identify up to four of the same type of filtering servers per context. In single mode a maximum of 16 of the same type of filtering 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

The URL Filtering Server Type area includes the following fields:

- 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.

The URL Filtering Servers area includes the following fields:

- 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. See the following topics for more information:

  - Add/Edit Parameters for Websense URL Filtering, page 27-3
  - Add/Edit Parameters for Secure Computing SmartFilter URL Filtering, page 27-4
- 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.

You can perform the following actions in this pane:

- 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

Advanced URL Filtering, page 27-4

Filter Rules, page 27-5

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

- 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

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. Choose 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. Choose 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. Valid values are from 1 to 128.
- 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

The Filter Rules pane 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 Configuration > Firewall > URL Filtering Servers pane. For more information, see URL Filtering, page 27-1.

Benefits

The Filter Rules pane 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 pane.

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 types of filter rules you can add. Clicking the rule type opens the Add Filter Rule dialog box for the specified filter rule type.
  - Add Filter ActiveX Rule
  - Add Filter Java Rule
  - Add Filter HTTP Rule
  - Add Filter HTTPS Rule
  - Add Filter FTP Rule
- Edit—Displays the Edit Filter Rule dialog box for editing the selected filtering rule.
- Delete—Deletes the selected filtering rule.
- 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 this button brings up an extended toolbar. See Filtering the Rule Table, page 27-9 for more information.
- Rule Diagram—Toggles the display of the Rule Diagram.
- Packet Trace—Launches the Packet Tracer utility.
- Use the Addresses tab to choose the source of the filter rule that you are choosing.
  - Type—Lets you choose a source from the drop-down list, 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 choose a predefined filter rule.
  - Type—Lets you choose a source from the drop-down list, 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 choose 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
•	•	•	•	—

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 window. For more information, see URL Filtering.

Fields

- Action—Provides the following drop-down list of different filtering actions to apply (the actions displayed depend upon the type of filter rule being created or edited):
  - 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
- Source—Enter the source of the traffic to which the filtering action applies. You can enter the source in one of the following ways:
  - any—Enter "any" (without quotation marks) to indicate any source address.
  - *name*—Enter a hostname.
  - *address*/*mask*—Enter an IP address and optional network mask. You can express the netmask in CIDR or dotted decimal notation. For example, you can enter 10.1.1.0/24 or 10.1.1.0/255.255.255.0.
  - ...—Opens the Browse Source dialog box. You can choose a host or address from the drop-down list.
- Destination—Identifies the destination of the traffic to which the filtering action applies. You can enter the destination in one of the following ways:
  - any—Enter "any" (without quotation marks) to indicate any destination address.
  - *name*—Enter a hostname.
  - *address*/*mask*—Enter an IP address and optional network mask. You can express the netmask in CIDR or dotted decimal notation. For example, you can enter 10.1.1.0/24 or 10.1.1.0/255.255.255.0.
  - ...—Opens the Browse Destination dialog box. You can choose a host or address from the drop-down list.
- Service —Identifies the service of the traffic to which the filtering action applies. You can enter the destination in one of the following ways:
  - tcp/*port*—The port number can be from 1 to 65535. Additionally, you can use the following modifiers with the TCP service:

    !=—Not equal to. For example, !=tcp/443

    <—Less than. For example, <tcp/2000.

    >—Great than. For example, >tcp/2000.

    - —Range. For example, tcp/2000-3000.
  - *name*—Enter a well-known service name, such as http or ftp.
  - ...—Opens the Browse Service dialog box. You can choose a service from the drop-down list.
- HTTP Options—This area appears only for HTTP filter rules.
  - When URL exceeds maximum permitted size—Choose the action to take when the URL exceeds the specified size. You can choose to truncate the URL or block the traffic.
  - 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 users from connecting to an HTTP proxy server—Prevent HTTP requests made through a proxy server.
  - **Truncate CGI parameters from URL sent to URL server**—The security appliance forwards only the CGI script location and the script name, without any parameters, to the filtering server.
- **HTTPS Options**—This area appears only when you choose the **Filter HTTPS** option from the drop-down list.
  - **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 Options**—This area appears only when you choose the **Filter FTP** option from the drop-down list.
  - **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 interactive FTP sessions (block if absolute FTP path is not provided)**—When enabled, FTP requests are dropped if they use a relative pathname 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
•	•	•	•	—

Filtering the Rule Table

It can be difficult to find a specific rule if your rule table includes a lot of entries. You can apply a filter to the rule table to show only the rules specified by the filter. To filter the rule table, perform the following steps:

Step 1 Click **Find** on the toolbar. The Filter toolbar appears.

Step 2 Choose the type of filter from the filter drop-down list:

- Source—Displays rules based on the specified source address or hostname.
- Destination—Displays rules based on the specified destination address or hostname.
- Source or Destination—Displays rules based on the specified source or destination address or hostname.
- Service—Displays rules based on the specified service.
- Rule Type—Displays rules based on the specified rule type.
- Query—Displays rules based on a complex query comprise of source, destination, service, and rule type information.

Step 3 For Source, Destination, Source or Destination, and Service filters, perform the following steps:

a. Choose the match criteria from the drop-down list. Choose "is" (without the quotes) for exact string matches or choose "contains" for partial string matches.

b. Enter the string to match using one of the following methods:

   – Type the source, destination, or service name into the condition field.

   – Click **...** to open a browse dialog from which you can choose existing services, IP addresses, or hostnames.

Step 4 For Rule Type filter, choose the rule type from the list.

Step 5 For Query filters, click **Define Query** and configure the complex query. For more information about configuring the complex query, see Browse Source/Destination/Service, page 27-11.

Step 6 To apply the filter to the rule table, click **Filter**.

Step 7 To clear the filter from the rule table and display all rule entries, click **Clear**.

Define Query

The Define Query dialog box lets you define a rule table filter based on multiple criteria, such as source, destination, service, and rule type.

Once you create the query and click OK, the filter is immediately applied to the rule table. You can clear the filter by clicking **Clear**.

Fields

- Source—IP address or hostname of the source. Choose "is" for an exact match or choose "contains" for a partial match. Click **...** to open up a selection dialog. You can specify a network mask using CIDR notation (address/bit-count). You can specify multiple addresses by separating them by commas (,).
- Destination—IP address or hostname of the destination. Choose "is" for an exact match or choose "contains" for a partial match. Click **...** to open up a selection dialog. You can specify a network mask using CIDR notation (address/bit-count). You can specify multiple addresses by separating them by commas (,).
- Source or Destination—IP address or hostname of the source or destination. Choose "is" for an exact match or choose "contains" for a partial match. Click **...** to open up a selection dialog. You can specify a network mask using CIDR notation (address/bit-count). You can specify multiple addresses by separating them by commas (,).
- Service—The protocol/port or name of a service. Choose "is" for an exact match or choose "contains" for a partial match. Click **...** to open up a selection dialog. You can specify multiple services by separating them by commas (,).
- Rule Type—Choose the rule type 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

Filtering the Rule Table, page 27-9

Browse Source/Destination/Service

The Browse Source/Destination/Service dialog box lets you choose from existing IP address, name, or service objects.

Fields

- Add—Click to add a new IP address, name, or service object.
- Edit—Click to edit an existing IP address, name, or service object.
- Filter/Clear—Enter a string by which to filter the information shown in the dialog box. Click Filter to apply the filter to the information shown in the dialog box. Click Clear to remove the filter and display all objects.
- Type—Organizes the objects shown into types, such as IP Names, IP Address Objects, and so on.
- Name—The name of the object. For services, it is the service name. For IP Address objects, it is the IP address, for IP name objects, it is the hostname.
- IP Address—The IP address of the address object.
- Netmask—The network mask of the address object.
- Protocol—The network protocol used by the service (such as tcp, udp, or icmp).
- Source Ports—The source port used by the service.
- Destination Ports—The destination port used by the service.
- ICMP Type—The ICMP type (for example 9, which is a router advertisement).
- Description (optional)—Specifies a description for the object.
- Source/Destination/Service button—Click this to add the address or service object to the filter rule or query.

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, page 27-5

URL Filtering, page 27-1

CHAPTER 28

Configuring Advanced Firewall Protection

This chapter describes how to prevent network attacks by configuring protection features, and includes the following sections:

Note For Sun RPC server and encrypted traffic inspection settings, which you configure in the Configuration > Firewall > Advanced area (along with many of the topics in this chapter), see Chapter 25, “Configuring Application Layer Protocol Inspection.”

Configuring Threat Detection

This section describes how to configure scanning threat detection and basic threat detection. Threat detection is available in single mode only.

This section includes the following topics:

To view threat detection statistics, see the Firewall Dashboard Tab.

Configuring Basic Threat Detection

Basic threat detection detects activity that might be related to an attack, such as a DoS attack. Basic threat detection is enabled by default.

This section includes the following topics:

Basic Threat Detection Overview

Using basic threat detection, the security appliance monitors the rate of dropped packets and security events due to the following reasons:

- Denial by access lists
- Bad packet format (such as invalid-ip-header or invalid-tcp-hdr-length)
- Connection limits exceeded (both system-wide resource limits, and limits set in the configuration)
- DoS attack detected (such as an invalid SPI, Stateful Firewall check failure)
- Basic firewall checks failed (This option is a combined rate that includes all firewall-related packet drops in this bulleted list. It does not include non-firewall-related drops such as interface overload, packets failed at application inspection, and scanning attack detected.)
- Suspicious ICMP packets detected
- Packets failed application inspection
- Interface overload
- Scanning attack detected (This option monitors scanning attacks; for example, the first TCP packet is not a SYN packet, or the TCP connection failed the 3-way handshake. Full scanning threat detection (see the "Configuring Scanning Threat Detection" section on page 28-3) takes this scanning attack rate information and acts on it by classifying hosts as attackers and automatically shunning them, for example.)
- Incomplete session detection such as TCP SYN attack detected or no data UDP session attack detected

When the security appliance detects a threat, it immediately sends a system log message (730100).

Basic threat detection affects performance only when there are drops or potential threats; even in this scenario, the performance impact is insignificant.

Configuring Basic Threat Detection

To enable or disable basic threat detection, on the Configuration > Firewall > Threat Detection pane, click the **Enable Basic Threat Detection** check box.

By default, this option enables detection for certain types of security events, including packet drops and incomplete session detections. You can override the default settings for each type of event if desired.

If an event rate is exceeded, then the security appliance sends a system message. The security appliance tracks two types of rates: the average event rate over an interval, and the burst event rate over a shorter burst interval. The burst rate interval is 1/30th of the average rate interval or 10 seconds, whichever is higher. For each received event, the security appliance checks the average and burst rate limits; if both rates are exceeded, then the security appliance sends two separate system messages, with a maximum of one message for each rate type per burst period.

Table 28-1 lists the default settings.

Table 28-1 Basic Threat Detection Default Settings

	Trigger Settings	
Packet Drop Reason	Average Rate	Burst Rate
• DoS attack detected • Bad packet format • Connection limits exceeded • Suspicious ICMP packets detected	100 drops/sec over the last 600 seconds.	400 drops/sec over the last 10 second period.
	80 drops/sec over the last 3600 seconds.	320 drops/sec over the last 60 second period.
Scanning attack detected	5 drops/sec over the last 600 seconds.	10 drops/sec over the last 10 second period.
	4 drops/sec over the last 3600 seconds.	8 drops/sec over the last 60 second period.
Incomplete session detected such as TCP SYN attack detected or no data UDP session attack detected (combined)	100 drops/sec over the last 600 seconds.	200 drops/sec over the last 10 second period.
	80 drops/sec over the last 3600 seconds.	160 drops/sec over the last 60 second period.
Denial by access lists	400 drops/sec over the last 600 seconds.	800 drops/sec over the last 10 second period.
	320 drops/sec over the last 3600 seconds.	640 drops/sec over the last 60 second period.
• Basic firewall checks failed • Packets failed application inspection	400 drops/sec over the last 600 seconds.	1600 drops/sec over the last 10 second period.
	320 drops/sec over the last 3600 seconds.	1280 drops/sec over the last 60 second period.
Interface overload	2000 drops/sec over the last 600 seconds.	8000 drops/sec over the last 10 second period.
	1600 drops/sec over the last 3600 seconds.	6400 drops/sec over the last 60 second period.

Configuring Scanning Threat Detection

A typical scanning attack consists of a host that tests the accessibility of every IP address in a subnet (by scanning through many hosts in the subnet or sweeping through many ports in a host or subnet). The scanning threat detection feature determines when a host is performing a scan. Unlike IPS scan detection that is based on traffic signatures, the security appliance scanning threat detection feature maintains an extensive database that contains host statistics that can be analyzed for scanning activity.

The host database tracks suspicious activity such as connections with no return activity, access of closed service ports, vulnerable TCP behaviors such as non-random IPID, and many more behaviors.

You can configure the security appliance to send system log messages about an attacker or you can automatically shun the host.

⚠ **Caution** The scanning threat detection feature can affect the security appliance performance and memory significantly while it creates and gathers host- and subnet-based data structure and information.

To configure scanning threat detection, perform the following steps:

Step 1 To enable scanning threat detection, on the Configuration > Firewall > Threat Detection pane, click the **Enable Scanning Threat Detection** check box.

By default, the system log message 730101 is generated when a host is identified as an attacker.

The security appliance identifies a host as an attacker or as a target if the scanning threat rate is exceeded. The security appliance tracks two types of rates: the average event rate over an interval, and the burst event rate over a shorter burst interval. The burst event rate is 1/30th of the average rate interval or 10 seconds, whichever is higher. For each event detected that is considered to be part of a scanning attack, the security appliance checks the average and burst rate limits. If either rate is exceeded for traffic sent from a host, then that host is considered to be an attacker. If either rate is exceeded for traffic received by a host, then that host is considered to be a target.

Table 28-2 lists the default rate limits for scanning threat detection.

Table 28-2 Default Rate Limits for Scanning Threat Detection

Average Rate	Burst Rate
5 drops/sec over the last 600 seconds.	10 drops/sec over the last 10 second period.
5 drops/sec over the last 3600 seconds.	10 drops/sec over the last 60 second period.

Step 2 (Optional) To automatically terminate a host connection when the security appliance identifies the host as an attacker, check the **Shun Hosts detected by scanning threat** check box.

Step 3 (Optional) To except host IP addresses from being shunned, enter an address in the **Networks excluded from shun** field.

You can enter multiple addresses or subnets separated by commas. To choose a network from the list of IP address objects, click the **...** button.

Step 4 (Optional) To set the duration of a shun for an attacking host, check **Set Shun Duration** and enter a value between 10 and 2592000 seconds. The default length is 3600 seconds (1 hour). To restore the default value, click **Set Default**.

Configuring Threat Statistics

You can configure the security appliance to collect extensive statistics. Threat detection statistics show both allowed and dropped traffic rates. By default, statistics for access lists are enabled.

To view threat detection statistics, see the Firewall Dashboard Tab.

Caution Enabling statistics can affect the security appliance performance, depending on the type of statistics enabled. Enabling statistics for hosts affects performance in a significant way; if you have a high traffic load, you might consider enabling this type of statistics temporarily. Enabling statistics for ports, however, has modest impact.

- To enable *all* statistics, in the Configuration > Firewall > Threat Detection > Scanning Threat Statistics area, click the **Enable All Statistics** radio button.
- To disable *all* statistics, on the Configuration > Firewall > Threat Detection pane, click the **Disable All Statistics** radio button.
- To enable only certain statistics, on the Configuration > Firewall > Threat Detection > Scanning Threat Statistics area, click the **Enable Only Following Statistics** radio button, end then check one or more of the following check boxes:
  - **Hosts**—Enables host statistics. The host statistics accumulate for as long as the host is active and in the scanning threat host database. The host is deleted from the database (and the statistics cleared) after 10 minutes of inactivity.
  - **Access Rules** (enabled by default)—Enables statistics for access rules.
  - **Port**—Enables statistics for TCP and UDP ports.
  - **Protocol**—Enables statistics for non-TCP/UDP IP protocols.
  - **TCP-Intercept**—Enables statistics for attacks intercepted by TCP Intercept (see the "Enabling Connection Limits, TCP Normalization , and TCP State Bypass" section on page 28-10 to enable TCP Intercept). After you check the TCP-Intercept option, you can set the following options in the Configuration > Firewall > Threat Detection > TCP Intercept Threat Detection area:

    Monitoring Window Size—Sets the size of the history monitoring window, between 1 and 1440 minutes. The default is 30 minutes. The security appliance samples the number of attacks 30 times during the rate interval, so for the default 30 minute period, statistics are collected every 60 seconds.

    Burst Threshold Rate—Sets the threshold for syslog message generation, between 25 and 2147483647. The default is 400 per second. When the burst rate is exceeded, syslog message 733104 is generated.

    Average Threshold Rate—Sets the average rate threshold for syslog message generation, between 25 and 2147483647. The default is 200 per second. When the average rate is exceeded, syslog message 733105 is generated.

    Click the **Set Default** button to restore the default values.
- To set the number of rate intervals maintained for host statistics, on the Configuration > Firewall > Threat Detection > Scanning Threat Statistics area, choose **1**, **2**, or **3** from the User can specify the number of rate for Threat Detection Host drop-down list. Because host statistics use a lot of memory, reducing the number of rate intervals from the default of 3 reduces the memory usage. By default, the Firewall Dashboard Tab shows information for three rate intervals, for example, for the last 1 hour, 8 hours, and 24 hours. If you set this keyword to **1**, then only the shortest rate interval statistics are maintained. If you set the value to **2**, then the two shortest intervals are maintained.

Configuring Connection Settings

This section describes how to:

- Set maximum TCP and UDP connections
- Set maximum embryonic connections
- Set maximum per-client connections
- Set connection timeouts
- Set dead connection detection
- Disable TCP sequence randomization.
- Configure TCP normalization.
- Configure TCP state bypass.

This section includes the following topics:

- Information About Connection Limits, page 28-6
- Information About TCP Normalization, page 28-7
- Information About TCP State Bypass, page 28-8
- Enabling Connection Limits, TCP Normalization , and TCP State Bypass, page 28-10

Note You can also configure maximum connections, maximum embryonic connections, and TCP sequence randomization in the NAT configuration. 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.

Information About Connection Limits

This section describes why you might want to limit connections, and includes the following topics:

- Information About TCP Intercept, page 28-6
- Disabling TCP Intercept for Management Packets for Clientless SSL VPN Compatibility, page 28-7
- Information About Dead Connection Detection, page 28-7
- Information About TCP Sequence Randomization, page 28-7

Information About TCP Intercept

Limiting the number of embryonic connections protects you from a DoS attack. The security appliance uses the per-client limits and the embryonic connection limit to trigger TCP Intercept, which protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. TCP Intercept uses the SYN cookies algorithm to prevent TCP SYN-flooding attacks. A SYN-flooding attack consists of a series of SYN packets usually originating from spoofed IP addresses. The constant flood of SYN packets keeps the server SYN queue full, which prevents it from servicing connection requests. When the embryonic connection threshold of a connection is crossed, the security appliance acts as a proxy for the server and generates a SYN-ACK response to the client SYN request. When the security appliance receives an ACK back from the client, it can then authenticate the client and allow the connection to the server.

Disabling TCP Intercept for Management Packets for Clientless SSL VPN Compatibility

By default, TCP management connections have TCP Intercept always enabled. When TCP Intercept is enabled, it intercepts the 3-way TCP connection establishment handshake packets and thus deprives the security appliance from processing the packets for Clientless (browser-based) SSL VPN. Clientless SSL VPN requires the ability to process the 3-way handshake packets to provide selective ACK and other TCP options for Clientless SSL VPN connections. To disable TCP Intercept for management traffic, you can set the embryonic connection limit; only after the embryonic connection limit is reached is TCP Intercept enabled.

Information About Dead Connection Detection

Dead connection detection detects a dead connection and allows it to expire, without expiring connections that can still handle traffic. You configure DCD when you want idle, but valid connections to persist.

When you enable DCD, idle timeout behavior changes. With idle timeout, DCD probes are sent to each of the two end-hosts to determine the validity of the connection. If an end-host fails to respond after probes are sent at the configured intervals, the connection is freed, and reset values, if configured, are sent to each of the end-hosts. If both end-hosts response that the connection is valid, the activity timeout is updated to the current time and the idle timeout is rescheduled accordingly.

Information About TCP Sequence Randomization

Each TCP connection has two ISNs: one generated by the client and one generated by the server. The security appliance randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

Randomizing the ISN of the protected host prevents an attacker from predecting the next ISN for a new connection and potentially hijacking the new session.

TCP initial sequence number randomization can be disabled if required. For example:

- If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.
- If you use eBGP multi-hop through the security appliance, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.
- You use a WAAS device that requires the security appliance not to randomize the sequence numbers of connections.

Information About TCP Normalization

The TCP normalization feature identifies abnormal packets that the security appliance can act on when they are detected; for example, the security appliance can allow, drop, or clear the packets. TCP normalization helps protect the security appliance from attacks. TCP normalization is always enabled, but you can customize how some features behave.

The TCP normalizer includes non-configurable actions and configurable actions. Typically, non-configurable actions that drop or clear connections apply to packets that are always bad. Configurable actions (as detailed in "Enabling Connection Limits, TCP Normalization , and TCP State Bypass" section on page 28-10) might need to be customized depending on your network needs.

See the following guidelines for TCP normalization:

- The normalizer does not protect from SYN floods. The security appliance includes SYN flood protection in other ways.
- The normalizer always sees the SYN packet as the first packet in a flow unless the security appliance is in loose mode due to failover.

Information About TCP State Bypass

By default, all traffic that goes through the security appliance is inspected using the Adaptive Security Algorithm and is either allowed through or dropped based on the security policy. The security appliance maximizes the firewall performance by checking the state of each packet (is this a new connection or an established connection?) and assigning it to either the session management path (a new connection SYN packet), the fast path (an established connection), or the control plane path (advanced inspection). See the "Stateful Inspection Overview" section on page 3-17 for more detailed information about the stateful firewall.

TCP packets that match existing connections in the fast path can pass through the security appliance without rechecking every aspect of the security policy. This feature maximizes performance. However, the method of establishing the session in the fast path using the SYN packet, and the checks that occur in the fast path (such as TCP sequence number), can stand in the way of asymmetrical routing solutions: both the outbound and inbound flow of a connection must pass through the same security appliance.

For example, a new connection goes to security appliance 1. The SYN packet goes through the session management path, and an entry for the connection is added to the fast path table. If subsequent packets of this connection go through security appliance 1, then the packets will match the entry in the fast path, and are passed through. But if subsequent packets go to security appliance 2, where there was not a SYN packet that went through the session management path, then there is no entry in the fast path for the connection, and the packets are dropped. Figure 28-1 shows an asymmetric routing example where the outbound traffic goes through a different security appliance than the inbound traffic:

Figure 28-1 Asymmetric Routing

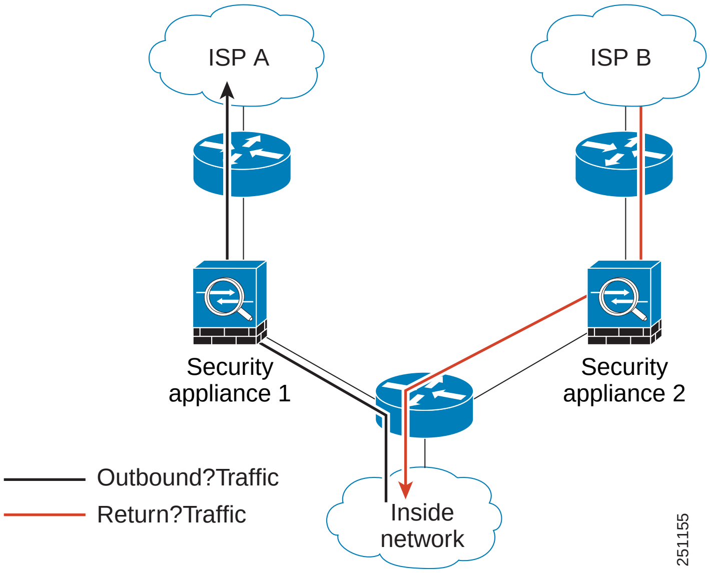

If you have asymmetric routing configured on upstream routers, and traffic alternates between two security appliances, then you can configure TCP state bypass for specific traffic. TCP state bypass alters the way sessions are established in the fast path and disables the fast path checks. This feature treats TCP traffic much as it treats a UDP connection: when a non-SYN packet matching the specified networks enters the security appliance, and there is not an fast path entry, then the packet goes through the session management path to establish the connection in the fast path. Once in the fast path, the traffic bypasses the fast path checks.

Licensing Requirements for TCP State Bypass

Model	License Requirement
All models	Base License.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent mode.

Failover Guidelines

Failover is supported.

IPv6 Guidelines

Supports IPv6.

Unsupported Features

The following features are not supported when you use TCP state bypass:

- Application inspection—Application inspection requires both inbound and outbound traffic to go through the same security appliance, so application inspection is not supported with TCP state bypass.
- AAA authenticated sessions—When a user authenticates with one security appliance, traffic returning via the other security appliance will be denied because the user did not authenticate with that security appliance.
- TCP Intercept, maximum embryonic connection limit, TCP sequence number randomization—The security appliance does not keep track of the state of the connection, so these features are not applied.
- TCP normalization—The TCP normalizer is disabled.
- SSM and SSC functionality—You cannot use TCP state bypass and any application running on an SSM or SSC, such as IPS or CSC.

NAT Guidelines

Because the translation session is established separately for each security appliance, be sure to configure static NAT on both security appliances for TCP state bypass traffic; if you use dynamic NAT, the address chosen for the session on security appliance 1 will differ from the address chosen for the session on security appliance 2.

Default Settings

TCP state bypass is disabled by default.

Enabling Connection Limits, TCP Normalization , and TCP State Bypass

To configure connection limits and TCP normalization, perform the following steps:

Step 1 Configure a service policy on the Configuration > Firewall > Service Policy Rules pane according to Chapter 23, “Configuring Service Policy Rules.”

You can configure connection limits as part of a new service policy rule, or you can edit an existing service policy.

Step 2 On the Rule Actions dialog box, click the **Connection Settings** tab.

Step 3 To set maximum connections, configure the following values in the Maximum Connections area:

- TCP & UDP Connections—Specifies the maximum number of simultaneous TCP and UDP connections for all clients in the traffic class, up to 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.

Step 4 To configure TCP timeouts, configure the following values in the 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.

Step 5 To disable randomized sequence numbers, uncheck **Randomize Sequence Number**.

TCP initial sequence number randomization can be disabled if another in-line firewall is also randomizing the initial sequence numbers, because there is no need for both firewalls to be performing this action. However, leaving ISN randomization enabled on both firewalls does not affect the traffic.

Each TCP connection has two ISNs: one generated by the client and one generated by the server. The security appliance randomizes the ISN of the TCP SYN passing in the outbound direction. If the connection is between two interfaces with the same security level, then the ISN will be randomized in the SYN in both directions.

Randomizing the ISN of the protected host prevents an attacker from predecting the next ISN for a new connection and potentially hijacking the new session.

Step 6 To configure TCP normalization, check **Use TCP Map**.

Choose an existing TCP map from the drop-down list (if available), or add a new one by clicking **New**.

The Add TCP Map dialog box appears.

a. In the TCP Map Name field, enter a name.

b. In the Queue Limit field, enter the maximum number of out-of-order packets, between 0 and 250 packets.

The Queue Limit sets the maximum number of out-of-order packets that can be buffered and put in order for a TCP connection. The default is 0, which means this setting is disabled and the default system queue limit is used depending on the type of traffic:

- Connections for application inspection, IPS, and TCP check-retransmission have a queue limit of 3 packets. If the security appliance receives a TCP packet with a different window size, then the queue limit is dynamically changed to match the advertised setting.
- For other TCP connections, out-of-order packets are passed through untouched.

If you set the Queue Limit command to be 1 or above, then the number of out-of-order packets allowed for all TCP traffic matches this setting. For application inspection, IPS, and TCP check-retransmission traffic, any advertised settings are ignored. For other TCP traffic, out-of-order packets are now buffered and put in order instead of passed through untouched.

c. In the Timeout field, set the maximum amount of time that out-of-order packets can remain in the buffer, between 1 and 20 seconds.

If they are not put in order and passed on within the timeout period, then they are dropped. The default is 4 seconds. You cannot change the timeout for any traffic if the Queue Limit is set to 0; you need to set the limit to be 1 or above for the Timeout to take effect.

d. In the Reserved Bits area, click **Clear and allow**, **Allow only**, or **Drop**.

Allow only allows packets with the reserved bits in the TCP header.

Clear and allow clears the reserved bits in the TCP header and allows the packet.

Drop drops the packet with the reserved bits in the TCP header.

e. Check any of the following options:

- Clear urgent flag—Clears the URG flag through the security appliance. The URG flag is used to indicate that the packet contains information that is of higher priority than other data within the stream. The TCP RFC is vague about the exact interpretation of the URG flag, therefore end systems handle urgent offsets in different ways, which may make the end system vulnerable to attacks.
- Drop connection on window variation—Drops a connection that has changed its window size unexpectedly. The window size mechanism allows TCP to advertise a large window and to subsequently advertise a much smaller window without having accepted too much data. From the TCP specification, “shrinking the window” is strongly discouraged. When this condition is detected, the connection can be dropped.
- Drop packets that exceed maximum segment size—Drops packets that exceed MSS set by peer.
- Check if transmitted data is the same as original—Enables the retransmit data checks.
- Drop packets which have past-window sequence—Drops packets that have past-window sequence numbers, namely the sequence number of a received TCP packet is greater than the right edge of the TCP receiving window. If you do not check this option, then the Queue Limit must be set to 0 (disabled).
- Drop SYN Packets with data—Drops SYN packets with data.
- Enable TTL Evasion Protection—Enables the TTL evasion protection offered by the security appliance. Do not enable this option if you want to prevent attacks that attempt to evade security policy.

  For example, an attacker can send a packet that passes policy with a very short TTL. When the TTL goes to zero, a router between the security appliance and the endpoint drops the packet. It is at this point that the attacker can send a malicious packet with a long TTL that appears to the security appliance to be a retransmission and is passed. To the endpoint host, however, it is the first packet that has been received by the attacker. In this case, an attacker is able to succeed without security preventing the attack.

- Verify TCP Checksum—Enables checksum verification.
- Drop SYNACK Packets with data—Drops TCP SYNACK packets that contain data.
- Drop packets with invalid ACK—Drops packets with an invalid ACK. You might see invalid ACKs in the following instances:
  - In the TCP connection SYN-ACK-received status, if the ACK number of a received TCP packet is not exactly same as the sequence number of the next TCP packet sending out, it is an invalid ACK.
  - Whenever the ACK number of a received TCP packet is greater than the sequence number of the next TCP packet sending out, it is an invalid ACK.

Note TCP packets with an invalid ACK are automatically allowed for WAAS connections.

f. To set TCP options, check any of the following options:

- 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.

g. Click **OK**.

Step 7 To set the time to live, check **Decrement time to live for a connection**.

Step 8 To enable TCP state bypass, in the Advanced Options area, check **TCP State Bypass**.

Step 9 Click **OK** or **Finish**.

Configuring 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

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

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

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 28-3 lists supported signatures and system message numbers.

Table 28-3 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 28-3 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 28-3 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 28-3 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	Informational	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 28-3 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	Informational	Triggers when a request is made to the portmapper for the YP update daemon (ypupdated) port.
6154	400046	ypxfrd (YP transfer daemon) Portmap Request	Informational	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.

Configuring the Fragment Size

By default, the security appliance allows up to 24 fragments per IP packet, and up to 200 fragments awaiting reassembly. You might need to let fragments on your network if you have an application that routinely fragments packets, such as NFS over UDP. However, if you do not have an application that fragments traffic, we recommend that you do not allow fragments through the security appliance. Fragmented packets are often used as DoS attacks.

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

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

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
•	•	•	•	—

Configuring 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
•	—	•	•	—

Configuring TCP Options

The TCP Options pane lets you set parameters for TCP connections.

Fields

- Inbound and Outbound Reset—Sets whether to reset denied TCP connections for inbound and outbound traffic.
  - Interface—Shows the interface name.
  - Inbound Reset—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—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 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.
  - Edit—Sets the inbound and outbound reset settings for the interface.
- Other Options—Sets additional TCP options.
  - Send Reset Reply for Denied Outside TCP Packets—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—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—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—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

This dialog box sets the inbound and outbound reset settings for an interface.

Fields

- Send Reset Reply for Denied Inbound TCP Packets—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—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
•	•	•	•	—

Configuring Global Timeouts

The Timeouts pane 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 Provisional Media—Modifies the timeout value for SIP provisional media connections, between 0:1:0 and 1193:0:0. The default is 2 minutes.
- SIP Invite—Modifies the idle time after which pinholes for PROVISIONAL responses and media xlates will be closed. The minimum value is 0:1:0, the maximum value is 0:30:0. The default value is 0:03:00.

- 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.

Note When Authentication Absolute = 0, HTTPS authentication may not work. If a browser initiates multiple TCP connections to load a web page after HTTPS authentication, the first connection is permitted through, but subsequent connections trigger authentication. As a result, users are continuously presented with an authentication page, even after successful authentication. To work around this, set the authentication absolute timeout to 1 second. This workaround opens a 1-second window of opportunity that might allow non-authenticated users to go through the firewall if they are coming from the same source IP address.

- Authentication inactivity—Modifies the idle time until the authentication cache times out and users have to reauthenticate a new connection. This duration must be shorter than the Translation Slot value.
- Translation Slot—Modifies the idle time until a translation slot is freed. This duration must be at least 1 minute. The default is 3 hours. Enter 0:0:0 to disable timeout.

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

Configuring the Botnet Traffic Filter

Malware is malicious software that is installed on an unknowing host. Malware that attempts network activity such as sending private data (passwords, credit card numbers, key strokes, or proprietary data) can be detected by the Botnet Traffic Filter when the malware starts a connection to a known bad IP address. The Botnet Traffic Filter checks incoming and outgoing connections against a dynamic database of known bad domain names and IP addresses (the *blacklist*), and then logs or blocks any suspicious activity.

You can also supplement the Cisco dynamic database with blacklisted addresses of your choosing by adding them to a static blacklist; if the dynamic database includes blacklisted addresses that you think should not be blacklisted, you can manually enter them into a static *whitelist*. Whitelisted addresses still generate syslog messages, but because you are only targeting blacklist syslog messages, they are informational.

Note If you do not want to use the Cisco dynamic database at all, because of internal requirements, you can use the static blacklist alone if you can identify all the malware sites that you want to target.

This chapter describes how to configure the Botnet Traffic Filter, and includes the following sections:

- Information About the Botnet Traffic Filter, page 29-1
- Licensing Requirements for the Botnet Traffic Filter, page 29-5
- Guidelines and Limitations, page 29-5
- Default Settings, page 29-6
- Configuring the Botnet Traffic Filter, page 29-6
- Monitoring the Botnet Traffic Filter, page 29-13
- Where to Go Next, page 29-15
- Feature History for the Botnet Traffic Filter, page 29-15

Information About the Botnet Traffic Filter

This section includes information about the Botnet Traffic Filter, and includes the following topics:

- Botnet Traffic Filter Address Categories, page 29-2
- Botnet Traffic Filter Actions for Known Addresses, page 29-2
- Botnet Traffic Filter Databases, page 29-2

- How the Botnet Traffic Filter Works, page 29-4

Botnet Traffic Filter Address Categories

Addresses monitored by the Botnet Traffic Filter include:

- Known malware addresses—These addresses are on the blacklist identified by the dynamic database and the static blacklist.
- Known allowed addresses—These addresses are on the whitelist. The whitelist is useful when an address is blacklisted by the dynamic database and also identified by the static whitelist.
- Ambiguous addresses—These addresses are associated with multiple domain names, but not all of these domain names are on the blacklist. These addresses are on the *greylist*.
- Unlisted addresses—These addresses are unknown, and not included on any list.

Botnet Traffic Filter Actions for Known Addresses

You can configure the Botnet Traffic Filter to log suspicious activity, and you can optionally configure it to block suspicious traffic automatically.

Unlisted addresses do not generate any syslog messages, but addresses on the blacklist, whitelist, and greylist generate syslog messages differentiated by type. See the "Botnet Traffic Filter Syslog Messaging" section on page 29-13 for more information.

Botnet Traffic Filter Databases

The Botnet Traffic Filter uses two databases for known addresses. You can use both databases together, or you can disable use of the dynamic database and use the static database alone. This section includes the following topics:

- Information About the Dynamic Database, page 29-2
- Information About the Static Database, page 29-3
- Information About the DNS Reverse Lookup Cache and DNS Host Cache, page 29-3

Information About the Dynamic Database

The Botnet Traffic Filter can receive periodic updates for the dynamic database from the Cisco update server. This database lists thousands of known bad domain names and IP addresses.

The security appliance uses the dynamic database as follows:

1. When the domain name in a DNS reply matches a name in the dynamic database, the Botnet Traffic Filter adds the name and IP address to the *DNS reverse lookup cache*.
2. When the infected host starts a connection to the IP address of the malware site, then the security appliance sends a syslog message informing you of the suspicious activity and optionally drops the traffic if you configured the security appliance to do so.
3. In some cases, the IP address itself is supplied in the dynamic database, and the Botnet Traffic Filter logs or drops any traffic to that IP address without having to inspect DNS requests.

The database files are stored in running memory; they are not stored in Flash memory. If you need to delete the database, use the Configuration > Firewall > Botnet Traffic Filter > Botnet Database pane Purge Botnet Database button instead. Be sure to first disable use of the database by unchecking the **Use Botnet data dynamically downloaded from updater server** check box in the Configuration > Firewall > Botnet Traffic Filter > Botnet Database > Dynamic Database Configuration area.

Note To use the database, be sure to configure a domain name server for the security appliance so that it can access the URL.

To use the domain names in the dynamic database, you need to enable DNS packet inspection with Botnet Traffic Filter snooping; the security appliance looks inside the DNS packets for the domain name and associated IP address.

Information About the Static Database

You can manually enter domain names or IP addresses (host or subnet) that you want to tag as bad names in a blacklist. Static blacklist entries are always designated with a Very High threat level. You can also enter names or IP addresses in a whitelist, so that names or addresses that appear on both the *dynamic* blacklist and the whitelist are identified only as whitelist addresses in syslog messages and reports. Note that you see syslog messages for whitelisted addresses even if the address is not also in the dynamic blacklist.

When you add a domain name to the static database, the security appliance waits 1 minute, and then sends a DNS request for that domain name and adds the domain name/IP address pairing to the *DNS host cache*. (This action is a background process, and does not affect your ability to continue configuring the security appliance). We recommend also enabling DNS packet inspection with Botnet Traffic Filter snooping. The security appliance uses Botnet Traffic Filter snooping instead of the regular DNS lookup to resolve static blacklist domain names in the following circumstances:

- The security appliance DNS server is unavailable.
- A connection is initiated during the 1 minute waiting period before the security appliance sends the regular DNS request.

If DNS snooping is used, when an infected host sends a DNS request for a name on the static database, the security appliance looks inside the DNS packets for the domain name and associated IP address and adds the name and IP address to the DNS reverse lookup cache.

If you do not enable Botnet Traffic Filter snooping, and one of the above circumstances occurs, then that traffic will not be monitored by the Botnet Traffic Filter.

Information About the DNS Reverse Lookup Cache and DNS Host Cache

When you use the dynamic database with DNS snooping, entries are added to the DNS reverse lookup cache. If you use the static database, entries are added to the DNS host cache (see the "Information About the Static Database" section on page 29-3 about using the static database with DNS snooping and the DNS reverse lookup cache).

Entries in the DNS reverse lookup cache and the DNS host cache have a time to live (TTL) value provided by the DNS server. The largest TTL value allowed is 1 day (24 hours); if the DNS server provides a larger TTL, it is truncated to 1 day maximum.

For the DNS reverse lookup cache, after an entry times out, the security appliance renews the entry when an infected host initiates a connection to a known address, and DNS snooping occurs.

For the DNS host cache, after an entry times out, the security appliance periodically requests a refresh for the entry.

For the DNS host cache, the maximum number of blacklist entries and whitelist entries is 1000 each.

Table 29-1 lists the maximum number of entries in the DNS reverse lookup cache per model.

Table 29-1 DNS Reverse Lookup Cache Entries per Model

ASA Model	Maximum Entries
ASA 5505	5000
ASA 5510	10,000
ASA 5520	20,000
ASA 5540	40,000
ASA 5550	40,000
ASA 5580	100,000

How the Botnet Traffic Filter Works

Figure 29-1 shows how the Botnet Traffic Filter works with the dynamic database plus DNS inspection with Botnet Traffic Filter snooping.

Figure 29-1 How the Botnet Traffic Filter Works with the Dynamic Database

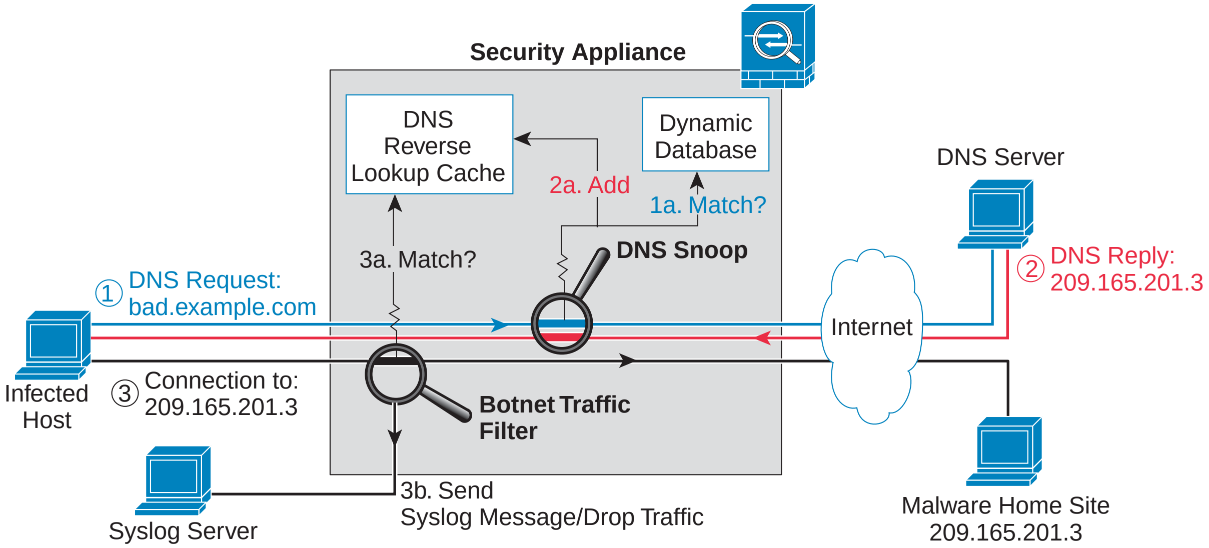

Figure 29-2 shows how the Botnet Traffic Filter works with the static database.

Figure 29-2 How the Botnet Traffic Filter Works with the Static Database

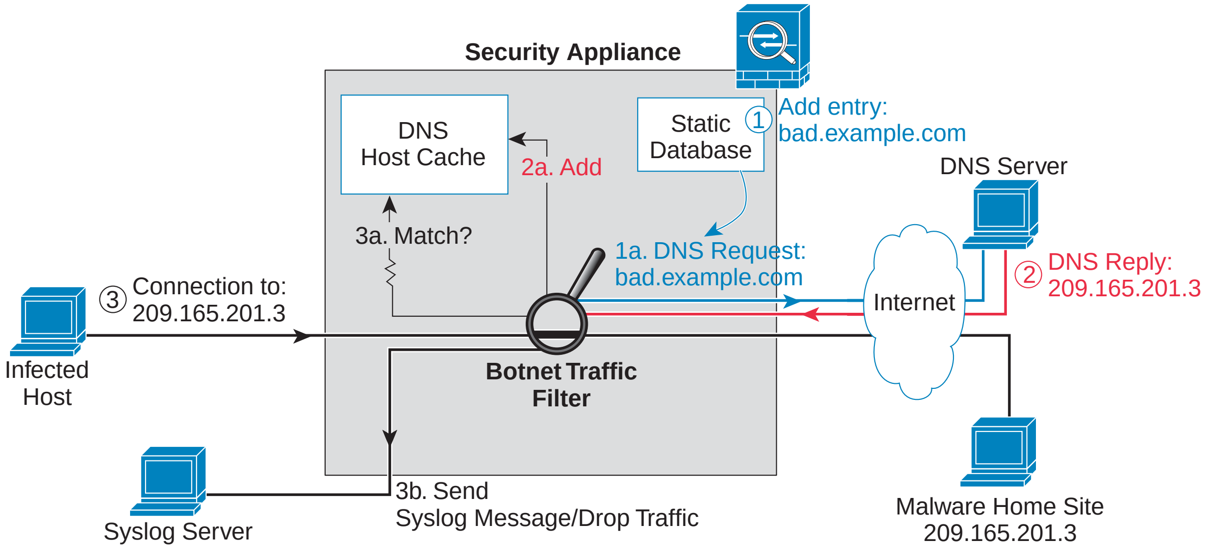

Licensing Requirements for the Botnet Traffic Filter

The following table shows the licensing requirements for this feature:

Model	License Requirement
All models	Botnet Traffic Filter License.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Failover Guidelines

Does not support replication of the DNS reverse lookup cache, DNS host cache, or the dynamic database in Stateful Failover.

IPv6 Guidelines

Does not support IPv6.

Additional Guidelines and Limitations

- TCP DNS traffic is not supported.
- You can add up to 1000 blacklist entries and 1000 whitelist entries in the static database.

Default Settings

By default, the Botnet Traffic Filter is disabled, as is use of the dynamic database.

For DNS inspection, which is enabled by default, Botnet Traffic Filter snooping is disabled by default.

Configuring the Botnet Traffic Filter

This section includes the following topics:

- Task Flow for Configuring the Botnet Traffic Filter, page 29-6
- Configuring the Dynamic Database, page 29-7
- Enabling DNS Snooping, page 29-9
- Adding Entries to the Static Database, page 29-8
- Enabling Traffic Classification and Actions for the Botnet Traffic Filter, page 29-10
- Blocking Botnet Traffic Manually, page 29-12
- Searching the Dynamic Database, page 29-13

Task Flow for Configuring the Botnet Traffic Filter

To configure the Botnet Traffic Filter, perform the following steps:

Step 1 Enable use of the dynamic database. See the "Configuring the Dynamic Database" section on page 29-7.

This procedure enables database updates from the Cisco update server, and also enables use of the downloaded dynamic database by the security appliance. Disallowing use of the downloaded database is useful in multiple context mode so you can configure use of the database on a per-context basis.

Step 2 (Optional) Add static entries to the database. See the "Adding Entries to the Static Database" section on page 29-8.

This procedure lets you augment the dynamic database with domain names or IP addresses that you want to blacklist or whitelist. You might want to use the static database instead of the dynamic database if you do not want to download the dynamic database over the Internet.

Step 3 Enable DNS snooping. See the "Enabling DNS Snooping" section on page 29-9.

This procedure enables inspection of DNS packets, compares the domain name with those in the dynamic database or the static database (when a DNS server for the security appliance is unavailable), and adds the name and IP address to the DNS reverse lookup cache. This cache is then used by the Botnet Traffic Filter when connections are made to the suspicious address.

Step 4 Enable traffic classification and actions for the Botnet Traffic Filter. See the "Enabling Traffic Classification and Actions for the Botnet Traffic Filter" section on page 29-10.

This procedure enables the Botnet Traffic Filter, which compares the source and destination IP address in each initial connection packet to the IP addresses in the dynamic database, static database, DNS reverse lookup cache, and DNS host cache, and sends a syslog message or drops any matching traffic.

Step 5 (Optional) Block traffic manually based on syslog message information. See the "Blocking Botnet Traffic Manually" section on page 29-12.

If you choose not to block malware traffic automatically, you can block traffic manually by configuring an access rule to deny traffic, or by using the **shun** command in the Command Line Interface tool to block all traffic to and from a host.

Configuring the Dynamic Database

This procedure enables database updates, and also enables use of the downloaded dynamic database by the security appliance. Disabling use of the downloaded database is useful in multiple context mode so you can configure use of the database on a per-context basis.

By default, downloading and using the dynamic database is disabled.

Prerequisites

Enable security appliance use of a DNS server in the Device Management > DNS > DNS Client > DNS Lookup area. In multiple context mode, enable DNS per context.

Detailed Steps

Step 1 Enable downloading of the dynamic database.

- In Single mode, choose the **Configuration > Firewall > Botnet Traffic Filter > Botnet Database** pane, then check the **Enable Botnet Updater Client** check box.
- In multiple context mode in the System execution space, choose the **Configuration > Device Management > Botnet Database** pane, then check the **Enable Botnet Updater Client** check box.

This setting enables downloading of the dynamic database from the Cisco update server. In multiple context mode, enter this command in the system execution space. If you do not have a database already installed on the security appliance, it downloads the database after approximately 2 minutes. The update server determines how often the security appliance polls the server for future updates, typically every hour.

Step 2 (Multiple context mode only) In multiple context mode, click **Apply**. Then change to the context where you want to configure the Botnet Traffic Filter by double-clicking the context name in the Device List.

Step 3 In the Configuration > Firewall > Botnet Traffic Filter > Botnet Database > Dynamic Database Configuration area, check the **Use Botnet data dynamically downloaded from updater server** check box.

Step 4 Click **Apply**.

Step 5 (Optional) If you want to later remove the database from running memory, perform the following steps:

a. Disable use of the database by unchecking the **Use Botnet data dynamically downloaded from updater server** check box.

b. Click **Apply**.

c. Click **Purge Botnet Database**.

d. To redownload the database, re-check the **Use Botnet data dynamically downloaded from updater server** check box.

e. Click **Apply**.

Note The Fetch Botnet Database button is for testing purposes only; it downloads and verifies the dynamic database, but does not store it in running memory.

For information about the Search Dynamic Database area, see the "Searching the Dynamic Database" section on page 29-13.

What to Do Next

See the "Adding Entries to the Static Database" section on page 29-8.

Adding Entries to the Static Database

The static database lets you augment the dynamic database with domain names or IP addresses that you want to blacklist or whitelist. Static blacklist entries are always designated with a Very High threat level. See the "Information About the Static Database" section on page 29-3 for more information.

Prerequisites

- In multiple context mode, perform this procedure in the context execution space.
- Enable security appliance use of a DNS server in the Device Management > DNS > DNS Client > DNS Lookup area. In multiple context mode, enable DNS per context.

Detailed Steps

Step 1 Choose the **Configuration** > **Firewall** > **Botnet Traffic Filter** > **Black** or White List pane, click **Add** for the Whitelist or Blacklist.

The Enter hostname or IP Address dialog box appears.

Step 2 In the Addresses field, enter one or more domain names, IP addresses, and IP address/netmasks.

Enter multiple entries separated by commas, spaces, lines, or semi-colons. You can enter up to 1000 entries for each type.

Step 3 Click **OK**.

Step 4 Click **Apply**.

What to Do Next

See the "Enabling DNS Snooping" section on page 29-9.

Enabling DNS Snooping

This procedure enables inspection of DNS packets and enables Botnet Traffic Filter snooping, which compares the domain name with those on the dynamic database or static database, and adds the name and IP address to the Botnet Traffic Filter DNS reverse lookup cache. This cache is then used by the Botnet Traffic Filter when connections are made to the suspicious address.

Prerequisites

- In multiple context mode, perform this procedure in the context execution space.
- You must first configure DNS inspection for traffic that you want to snoop using the Botnet Traffic Filter. See the "DNS Inspection" section on page 25-6 and Chapter 23, "Configuring a Service Policy Rule for Management Traffic," for detailed information about configuring advanced DNS inspection options using the Modular Policy Framework.

Note You can also configure DNS snooping directly in the Configuration > Firewall > Service Policy Rules > Rule Actions > Protocol Inspection > Select DNS Inspect Map dialog box by checking the **Enable Botnet traffic filter DNS snooping** check box.

Restrictions

TCP DNS traffic is not supported.

Default DNS Inspection Configuration and Recommended Configuration

The default configuration for DNS inspection inspects all UDP DNS traffic on all interfaces, and does not have DNS snooping enabled.

We suggest that you enable DNS snooping only on interfaces where external DNS requests are going. Enabling DNS snooping on all UDP DNS traffic, including that going to an internal DNS server, creates unnecessary load on the security appliance.

For example, if the DNS server is on the outside interface, you should enable DNS inspection with snooping for all UDP DNS traffic on the outside interface.

Detailed Steps

Step 1 Choose the **Configuration** > **Firewall** > **Botnet Traffic Filter** > **DNS Snooping** pane.

All existing service rules that include DNS inspection are listed in the table.

Step 2 For each rule for which you want to enable DNS snooping, in the DNS Snooping Enabled column, check the check box.

Step 3 Click **Apply**.

What to Do Next

See the "Enabling Traffic Classification and Actions for the Botnet Traffic Filter" section on page 29-10.

Enabling Traffic Classification and Actions for the Botnet Traffic Filter

This procedure enables the Botnet Traffic Filter. The Botnet Traffic Filter compares the source and destination IP address in each initial connection packet to the following:

- Dynamic database IP addresses
- Static database IP addresses
- DNS reverse lookup cache (for dynamic database domain names)
- DNS host cache (for static database domain names)

When an address matches, the security appliance sends a syslog message. The only additional action currently available is to drop the connection.

Prerequisites

In multiple context mode, perform this procedure in the context execution space.

Recommended Configuration

Although DNS snooping is not required, we recommend configuring DNS snooping for maximum use of the Botnet Traffic Filter (see the "Enabling DNS Snooping" section on page 29-9). Without DNS snooping for the dynamic database, the Botnet Traffic Filter uses only the static database entries, plus any IP addresses in the dynamic database; domain names in the dynamic database are not used.

We recommend enabling the Botnet Traffic Filter on all traffic on the Internet-facing interface, and enabling dropping of traffic with a severity of moderate and higher.

Detailed Steps

Step 1 Choose the **Configuration** > **Firewall** > **Botnet Traffic Filter** > **Traffic Settings** pane.

Step 2 To enable the Botnet Traffic Filter on specified traffic, perform the following steps:

a. In the Traffic Classification area, check the **Traffic Classified** check box for each interface on which you want to enable the Botnet Traffic Filter.

You can configure a global classification that applies to all interfaces by checking the Traffic Classified check box for Global (All Interfaces). If you configure an interface-specific classification, the settings for that interface overrides the global setting.

b. For each interface, from the **ACL Used** drop-down list choose either --ALL TRAFFIC-- (the default), or any access list configured on the security appliance.

For example, you might want to monitor all port 80 traffic on the outside interface.

To add or edit access lists, click **Manage ACL** to bring up the ACL Manager. See the "ACL Manager" section on page 36-14 for more information.

Step 3 (Optional) To treat greylisted traffic as blacklisted traffic for action purposes, in the Ambiguous Traffic Handling area, check the **Treat ambiguous (greylisted) traffic as malicious (blacklisted) traffic** check box.

If you do not enable this option, greylisted traffic will not be dropped if you configure a rule in the Blacklisted Traffic Actions area. See the "Botnet Traffic Filter Address Categories" section on page 29-2 for more information about the greylist.

Step 4 (Optional) To automatically drop malware traffic, perform the following steps.

To manually drop traffic, see the "Blocking Botnet Traffic Manually" section on page 29-12.

a. In the Blacklisted Traffic Actions area, click **Add**.

The Add Blacklisted Traffic Action dialog box appears.

b. From the Interface drop-down list, choose the interface on which you want to drop traffic. Only interfaces on which you enabled Botnet Traffic Filter traffic classification are available.

c. In the Threat Level area, choose one of the following options to drop traffic specific threat levels. The default level is a range between Moderate and Very High.

Note We highly recommend using the default setting unless you have strong reasons for changing the setting.

- Value—Specify the threat level you want to drop:
  - **Very Low**
  - **Low**
  - **Moderate**
  - **High**
  - **Very High**

Note Static blacklist entries are always designated with a Very High threat level.

- Range—Specify a range of threat levels.

d. In the ACL Used area, from the **ACL Used** drop-down list choose either --ALL TRAFFIC-- (the default), or any access list configured on the security appliance.

Note Be sure the access list is a subset of the traffic you specified in the Traffic Classification area.

To add or edit access lists, click **Manage** to bring up the ACL Manager. See the "ACL Manager" section on page 36-14 for more information.

e. Click **OK**.

You return to the Traffic Settings pane.

f. If you want to apply additional rules to a given interface, repeat steps a through e.

Make sure you do not specify overlapping traffic in multiple rules for a given interface. Because you cannot control the exact order that rules are matched, overlapping traffic means you do not know which command will be matched. For example, do not specify both a rule that matches --ALL TRAFFIC-- as well as a command with and access list for a given interface. In this case, the traffic might never match the command with the access list. Similarly, if you specify multiple commands with access lists, make sure each access list is unique, and that the networks do not overlap.

Step 5 Click **Apply**.

Blocking Botnet Traffic Manually

If you choose not to block malware traffic automatically (see the "Enabling Traffic Classification and Actions for the Botnet Traffic Filter" section on page 29-10), you can block traffic manually by configuring an access rule to deny traffic, or by using the **shun** command in the Command Line Interface tool to block all traffic to and from a host. For some messages, you can automatically configure access rules in ASDM.

For example, you receive the following syslog message:

```
ASA-4-338002: Dynamic Filter permitted black listed TCP traffic from inside:10.1.1.45/6798
(209.165.201.1/7890) to outside:209.165.202.129/80 (209.165.202.129/80), destination
209.165.202.129 resolved from dynamic list: bad.example.com
```

You can then perform one of the following actions:

- Create an access rule to deny traffic.

  For example, using the syslog message above, you might want to deny traffic from the infected host at 10.1.1.45 to the malware site at 209.165.202.129. Or, if there are many connections to different blacklisted addresses, you can create an access list to deny all traffic from 10.1.1.45 until you resolve the infection on the host computer.

  For the following syslog messages, a reverse access rule can be automatically created from the Real Time Log Viewer:

  – 338001, 338002, 338003, 338004 (blacklist)

  – 338201, 338202 (greylist)

  See Chapter 19, "Configuring Logging," and Chapter 21, "Configuring Access Rules and ACLs," for more information about creating an access rule.

  Note If you create a reverse access rule form a Botnet Traffic Filter syslog message, and you do not have any other access rules applied to the interface, then you might inadvertently block all traffic. Normally, without an access rule, all traffic from a high security to a low security interface is allowed. But when you apply an access rule, all traffic is denied except traffic that you explicitly permit. Because the reverse access rule is a deny rule, be sure to edit the resulting access policy for the interface to permit other traffic.

  Access lists block all future connections. To block the current connection, if it is still active, enter the **clear conn** command. For example, to clear only the connection listed in the syslog message, enter the **clear conn address 10.1.1.45 address 209.165.202.129** command. See the *Cisco ASA 5500 Series Command Reference* for more information.

- Shun the infected host.

  Shunning blocks all connections from the host, so you should use an access list if you want to block connections to certain destination addresses and ports. To shun a host, enter the following command in Tools > Command Line Interface. To drop the current connection as well as blocking all future connections, enter the destination address, source port, destination port, and optional protocol.

  ```
  shun src_ip [dst_ip src_port dest_port [protocol]]
  ```

  For example, to block future connections from 10.1.1.45, and also drop the current connection to the malware site in the syslog message, enter:

  ```
  shun 10.1.1.45 209.165.202.129 6798 80
  ```

After you resolve the infection, be sure to remove the access list or the shun. To remove the shun, enter **no shun** *src_ip*.

Searching the Dynamic Database

If you want to check if a domain name or IP address is included in the dynamic database, you can search the database for a string.

Detailed Steps

Step 1 Go to the Search Dynamic Database area:

- In Single mode or within a context, choose the **Configuration** > **Firewall** > **Botnet Traffic Filter** > **Botnet Database Update** pane.
- In multiple context mode in the System execution space, choose the **Configuration** > **Device Management** > **Botnet Database Update** pane.

Step 2 In the Search string field, enter a string at least 3 characters in length, and click **Find Now**.

The first two matches are shown. To refine your search for a more specific match, enter a longer string.

Step 3 To clear the displayed matches and the search string, click **Clear**, or you can just enter a new string and click **Find Now** to get a new display.

Monitoring the Botnet Traffic Filter

Whenever a known address is classified by the Botnet Traffic Filter, then a syslog message is generated. You can also monitor Botnet Traffic Filter statistics and other parameters by entering commands on the security appliance. This section includes the following topics:

- Botnet Traffic Filter Syslog Messaging, page 29-13
- Botnet Traffic Filter Monitor Panes, page 29-14

Botnet Traffic Filter Syslog Messaging

The Botnet Traffic Filter generates detailed syslog messages numbered 338*nnn*. Messages differentiate between incoming and outgoing connections, blacklist, whitelist, or greylist addresses, and many other variables. (The greylist includes addresses that are associated with multiple domain names, but not all of these domain names are on the blacklist.)

See the *Cisco ASA 5500 Series System Log Messages* for detailed information about syslog messages.

For the following syslog messages, a reverse access rule can be automatically created from the Real Time Log Viewer:

- 338001, 338002, 338003, 338004 (blacklist)
- 338201, 338202 (greylist)

See Chapter 19, “Configuring Logging.”

Botnet Traffic Filter Monitor Panes

To monitor the Botnet Traffic Filter, see the following panes:

Command	Purpose
Home > Firewall Dashboard	Shows the Top Botnet Traffic Filter Hits, which shows reports of the top 10 malware sites, ports, and infected hosts. This report is a snapshot of the data, and may not match the top 10 items since the statistics started to be collected. If you right-click an IP address, you can invoke the whois tool to learn more about the botnet site. • Top Malware Sites—Shows top malware sites. • Top Malware Ports—Shows top malware ports. • Top Infected Hosts—Shows the top infected hosts.
Monitoring > Botnet Traffic Filter > Statistics	Shows how many connections were monitored and dropped with the Botnet Traffic Filter, and how many of those connections match the whitelist, blacklist, and greylist. (The greylist includes addresses that are associated with multiple domain names, but not all of these domain names are on the blacklist.) The Details button shows how many packets at each threat level were classified or dropped.
Monitoring > Botnet Traffic Filter > Real-time Reports	Generates reports of the top 10 malware sites, ports, and infected hosts monitored. The top 10 malware-sites report includes the number of connections dropped, and the threat level and category of each site. This report is a snapshot of the data, and may not match the top 10 items since the statistics started to be collected. If you right-click a site IP address, you can invoke the whois tool to learn more about the malware site. Reports can be saved as a PDF file.
Monitoring > Botnet Traffic Filter > Infected Hosts	Generates reports about infected hosts. These reports contain detailed history about infected hosts, showing the correlation between infected hosts, visited malware sites, and malware ports. The Maximum Connections option shows the 20 infected hosts with the most number of connections. The Latest Activity option shows the 20 hosts with the most recent activity. The Highest Threat Level option shows the 20 hosts that connected to the malware sites with the highest threat level. The Subnet option shows up to 20 hosts within the specified subnet. Reports can be saved as a PDF file, as either the Current View or the Whole Buffer. The Whole Buffer option shows all buffered infected-hosts information.
Monitoring > Botnet Traffic Filter > Updater Client	Shows information about the updater server, including the server IP address, the next time the security appliance will connect with the server, and the database version last installed.
Monitoring > Botnet Traffic Filter > DNS Snooping	Shows the Botnet Traffic Filter DNS snooping actual IP addresses and names. All inspected DNS data is included in this output, and not just matching names in the blacklist. DNS data from static entries are not included.

Command	Purpose
Monitoring > Botnet Traffic Filter > Dynamic Database	Shows information about the dynamic database, including when the dynamic database was last downloaded, the version of the database, how many entries the database contains, and 10 sample entries.
Monitoring > Botnet Traffic Filter > ASP Table Hits	Shows the Botnet Traffic Filter rules that are installed in the accelerated security path.

Where to Go Next

- To configure the syslog server, see Chapter 19, "Configuring Logging."
- To block connections with an access rule, see Chapter 21, "Configuring Access Rules and ACLs."

Feature History for the Botnet Traffic Filter

Table 29-2 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

Table 29-2 Feature History for the Botnet Traffic Filter

Feature Name	Platform Releases	Feature Information
Botnet Traffic Filter	8.2(1)	This feature was introduced.
Automatic blocking, and blacklist category and threat level reporting.	8.2(2)	The Botnet Traffic Filter now supports automatic blocking of blacklisted traffic based on the threat level. You can also view the category and threat level of malware sites in statistics and reports. The following screens were introduced or modified: Configuration > Firewall > Botnet Traffic Filter > Traffic Settings, and Monitoring > Botnet Traffic Filter > Infected Hosts.

CHAPTER 30

Configuring the IPS Application on the AIP SSM and SSC

This chapter describes how to configure the IPS application that runs on an AIP SSM or an AIP SSC.

Note The SSC is supported on the ASA 5505. See the "SSM and SSC Support Per Model" section on page 3-2 for more information about which models support SSMs.

This chapter includes the following sections:

Information About the AIP SSM and SSC

You can install the AIP SSM/SSC into an ASA 5500 series adaptive security appliance. The AIP SSM/SSC runs advanced IPS software that provides proactive, full-featured intrusion prevention services to stop malicious traffic, including worms and network viruses, before they can affect your network. This section includes the following topics:

How the AIP SSM/SSC Works with the Adaptive Security Appliance

The AIP SSM/SSC runs a separate application from the security appliance. It is, however, integrated into the security appliance traffic flow. The AIP SSM/SSC does not contain any external interfaces itself (except for the management interface on the SSM only). When you identify traffic for IPS inspection on the security appliance, traffic flows through the security appliance and the AIP SSM/SSC in the following way:

1. Traffic enters the security appliance.
2. Firewall policies are applied.
3. Traffic is sent to the AIP SSM/SSC over the backplane.

   See the "Operating Modes" section on page 30-2 for information about only sending a copy of the traffic to the AIP SSM/SSC.
4. The AIP SSM/SSC applies its security policy to the traffic, and takes appropriate actions.
5. Valid traffic is sent back to the adaptive security appliance over the backplane in inline mode; the AIP SSM/SSC might block some traffic according to its security policy, and that traffic is not passed on.
6. VPN policies are applied (if configured).
7. Traffic exits the adaptive security appliance.

Figure 30-1 shows the traffic flow when running the AIP SSM/SSC in inline mode. In this example, the AIP SSM/SSC automatically blocks traffic that it identified as an attack. All other traffic is forwarded through the security appliance.

Figure 30-1 AIP SSM/SSC Traffic Flow in the Adaptive Security Appliance: Inline Mode

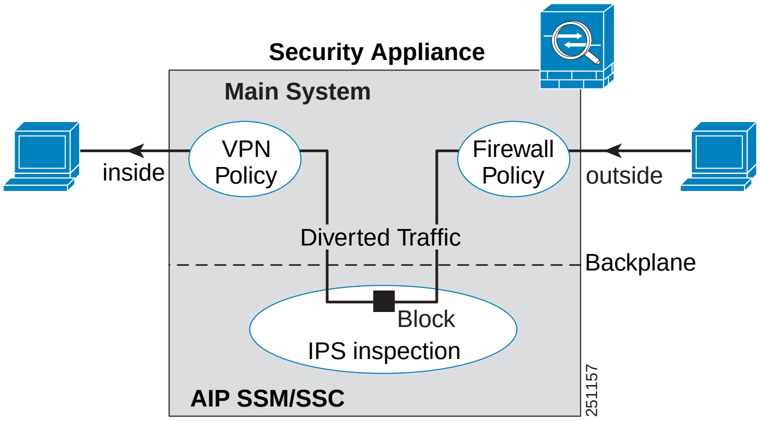

Operating Modes

You can send traffic to the AIP SSM/SSC using one of the following modes:

- Inline mode—This mode places the AIP SSM/SSC directly in the traffic flow Figure 30-1. No traffic that you identified for IPS inspection can continue through the adaptive security appliance without first passing through, and being inspected by, the AIP SSM/SSC. This mode is the most secure because every packet that you identify for inspection is analyzed before being allowed through. Also, the AIP SSM/SSC can implement a blocking policy on a packet-by-packet basis. This mode, however, can affect throughput.

- Promiscuous mode—This mode sends a duplicate stream of traffic to the AIP SSM/SSC. This mode is less secure, but has little impact on traffic throughput. Unlike the inline mode, in promiscuous mode the AIP SSM/SSC can only block traffic by instructing the adaptive security appliance to shun the traffic or by resetting a connection on the adaptive security appliance. Also, while the AIP SSM/SSC is analyzing the traffic, a small amount of traffic might pass through the adaptive security appliance before the AIP SSM/SSC can shun it.
- Figure 30-2 shows the AIP SSM/SSC in promiscuous mode. In this example, the AIP SSM/SSC sends a shun message to the security appliance for traffic it identified as a threat.

Figure 30-2 AIP SSM/SSC Traffic Flow in the Adaptive Security Appliance: Promiscuous Mode

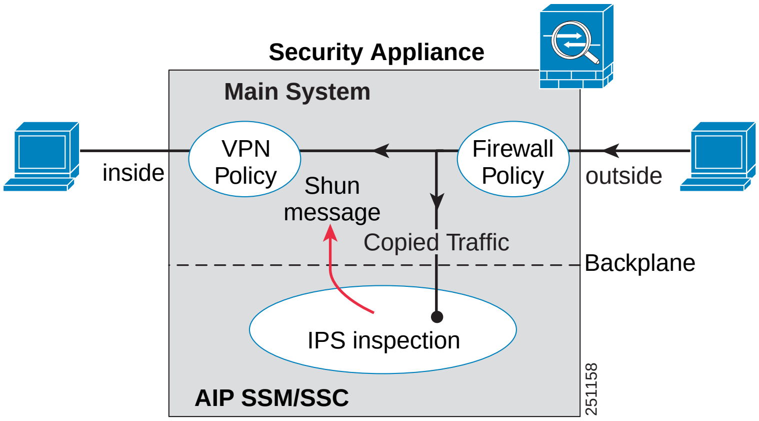

Using Virtual Sensors (AIP SSM Only)

The AIP SSM running IPS software Version 6.0 and above can run multiple virtual sensors, which means you can configure multiple security policies on the AIP SSM. You can assign each context or single mode security appliance to one or more virtual sensors, or you can assign multiple security contexts to the same virtual sensor. See the IPS documentation for more information about virtual sensors, including the maximum number of sensors supported.

Figure 30-3 shows one security context paired with one virtual sensor (in inline mode), while two security contexts share the same virtual sensor.

Figure 30-3 Security Contexts and Virtual Sensors

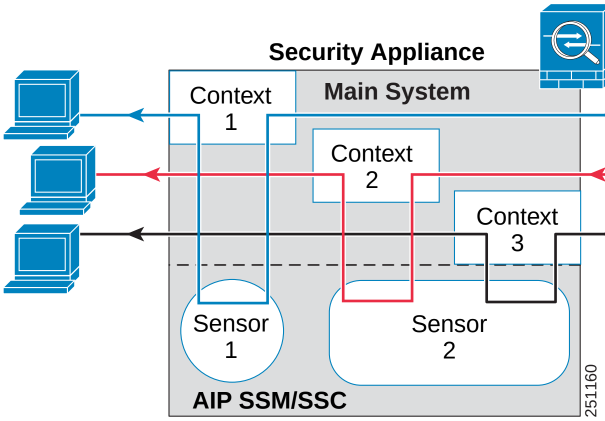

Figure 30-4 shows a single mode security appliance paired with multiple virtual sensors (in inline mode); each defined traffic flow goes to a different sensor.

Figure 30-4 Single Mode Security Appliance with Multiple Virtual Sensors

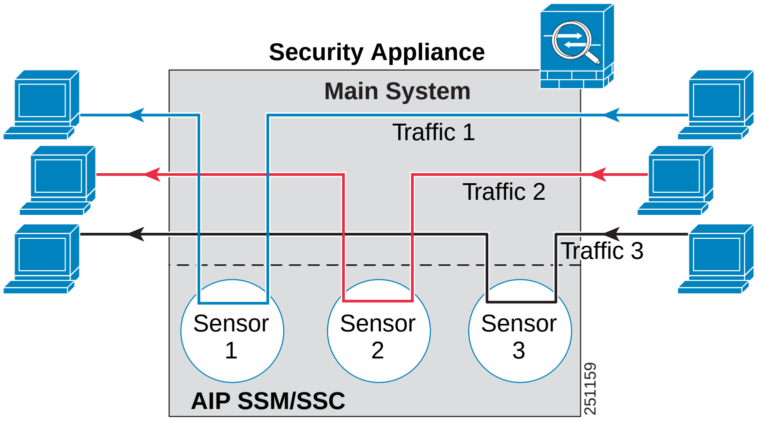

Differences Between the AIP SSM and the AIP SSC

The AIP SSM supports the higher performance requirements of the ASA 5510 and above, while the AIP SSC is designed for the small office installation of the ASA 5505 security appliance. The following features are supported on the AIP SSM, but not on the AIP SSC:

- Virtual sensors
- Anomaly detection
- Unretirement of default retired signatures

Licensing Requirements for the AIP SSM/SSC

The following table shows the licensing requirements for this feature:

Model	License Requirement
All models	Base License.

The IPS application on the AIP SSM/SSC requires a separate Cisco Services for IPS license in order to support signature updates. All other updates are available without a license.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

The ASA 5505 adaptive security appliance does not support multiple context mode, so multiple context features, such as virtual sensors, are not supported on the AIP SSC.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Model Guidelines

- The SSC is supported on the ASA 5505 only. See the "Supported Platforms and SSMs" section on page 2-2 for more information about which models support SSMs.
- The ASA 5505 adaptive security appliance does not support multiple context mode, so multiple context features, such as virtual sensors, are not supported on the AIP SSC.

Configuring the AIP SSM/SSC

This section describes how to configure IPS for the AIP SSM and AIP SSC, and includes the following topics:

- AIP SSM/SSC Task Overview, page 30-6
- Configuring the Security Policy on the AIP SSM/SSC, page 30-6
- Assigning Virtual Sensors to Security Contexts (AIP SSM Only), page 30-7
- Diverting Traffic to the AIP SSM/SSC, page 30-8
- Resetting the AIP SSM/SSC Password, page 30-9

AIP SSM/SSC Task Overview

Configuring the AIP SSM/SSC is a process that includes configuration of the IPS software on the SSM/SSC and then configuration of the ASA 5500 series adaptive security appliance. To configure the AIP SSM/SSC, perform the following steps:

Step 1 From ASDM, launch IDM. ASDM uses IDM to configure the AIP SSM. In IDM, configure the inspection and protection policy, which determines how to inspect traffic and what to do when an intrusion is detected. For the AIP SSM only, configure the inspection and protection policy for each virtual sensor if you want to run the AIP SSM in multiple sensor mode. See the "Configuring the Security Policy on the AIP SSM/SSC" section on page 30-6.

Step 2 (AIP SSM only) Using ASDM on the ASA 5500 in multiple context mode, specify which IPS virtual sensors are available for each context (if you configured virtual sensors). See the "Assigning Virtual Sensors to Security Contexts (AIP SSM Only)" section on page 30-7.

Step 3 Using ASDM on the ASA 5500, identify traffic to divert to the AIP SSM/SSC. See the "Diverting Traffic to the AIP SSM/SSC" section on page 30-8.

Configuring the Security Policy on the AIP SSM/SSC

ASDM uses IDM to configure the AIP SSM. This section describes how to access IDM from within ASDM.

Note See also the "Configuring the SSC Management Interface" section on page 10-3 to configure the SSC management interface for ASDM access and other uses.

Detailed Steps

Step 1 To access IDM from ASDM, click **Configuration** > **IPS**.

Step 2 You are asked for the IP address or hostname of the AIP SSM/SSC.

- If the AIP SSM/SSC is running IPS Version 6.0 or later, ASDM retrieves IDM from the AIP SSM/SSC and displays it as part of the ASDM interface. Enter the AIP SSM/SSC password and click **OK**.

  The IDM panes appear in the ASDM window.

- For the AIP SSM only, if it is running an earlier version of IPS software, ASDM displays a link to IDM. Click the link to launch IDM in a new browser window. You need to provide a username and password to access IDM.

If the password to access IDM is lost, you can reset the password using ASDM. See the "Resetting the AIP SSM/SSC Password" section on page 30-9, for more information.

Step 3 Configure the IPS security policy.

For the AIP SSM only, if you configure virtual sensors in IPS Version 6.0 or above, you identify one of the sensors as the default. If the ASA 5500 series adaptive security appliance does not specify a virtual sensor name in its configuration, the default sensor is used.

Because the IPS software that runs on the AIP SSM/SSC is beyond the scope of this document, detailed configuration information is available in the IPS documents at the following location:

http://www.cisco.com/en/US/products/hw/vpndevc/ps4077/tsd_products_support_series_home.html

What to Do Next

For the security appliance in multiple context mode, see the "Assigning Virtual Sensors to Security Contexts (AIP SSM Only)" section on page 30-7.

For the security appliance in single context mode, see the "Diverting Traffic to the AIP SSM/SSC" section on page 30-8.

Assigning Virtual Sensors to Security Contexts (AIP SSM Only)

If the security appliance is in multiple context mode, then you can assign one or more IPS virtual sensors to each context. Then, when you configure the context to send traffic to the AIP SSM, you can specify a sensor that is assigned to the context; you cannot specify a sensor that you did not assign to the context. If you do not assign any sensors to a context, then the default sensor configured on the AIP SSM is used. You can assign the same sensor to multiple contexts.

Note You do not need to be in multiple context mode to use virtual sensors; you can be in single mode and use different sensors for different traffic flows.

Detailed Steps

Step 1 In the ASDM Device List pane, double-click **System** under the active device IP address.

Step 2 On the Context Management > Security Contexts pane, choose a context that you want to configure, and click **Edit**.

The Edit Context dialog box appears. For more information about configuring contexts, see the "Configuring Security Contexts" section on page 11-16.

Step 3 In the IPS Sensor Allocation area, click **Add**.

The IPS Sensor Selection dialog box appears.

Step 4 From the Sensor Name drop-down list, choose a sensor name from those configured on the AIP SSM.

Step 5 (Optional) To assign a mapped name to the sensor, enter a value in the Mapped Sensor Name field.

This sensor name can be used within the context instead of the actual sensor name. If you do not specify a mapped name, the sensor name is used within the context. For security purposes, you might not want the context administrator to know which sensors are being used by the context. Or you might want to genericize the context configuration. For example, if you want all contexts to use sensors called "sensor1" and "sensor2," then you can map the "highsec" and "lowsec" sensors to sensor1 and sensor2 in context A, but map the "medsec" and "lowsec" sensors to sensor1 and sensor2 in context B.

Step 6 Click **OK** to return to the Edit Context dialog box.

Step 7 (Optional) To set one sensor as the default sensor for this context, from the Default Sensor drop-down list, choose a sensor name.

If you do not specify a sensor name when you configure IPS within the context configuration, the context uses this default sensor. You can only configure one default sensor per context. If you do not specify a sensor as the default, and the context configuration does not include a sensor name, then traffic uses the default sensor on the AIP SSM.

Step 8 Repeat this procedure for each security context.

Step 9 Change to each context to configure the IPS security policy as described in "Diverting Traffic to the AIP SSM/SSC" section on page 30-8.

What to Do Next

Change to each context to configure the IPS security policy as described in "Diverting Traffic to the AIP SSM/SSC" section on page 30-8.

Diverting Traffic to the AIP SSM/SSC

To identify traffic to divert from the adaptive security appliance to the AIP SSM/SSC, perform the following steps. In multiple context mode, perform these steps in each context execution space.

This feature is enabled using Service Policy rules. See Chapter 23, "Configuring Service Policy Rules," for detailed information about creating a service policy.

To configure the IPS service policy, perform the following steps:

Step 1 In the ASDM Device List pane, double-click the context name under the active device *IP address* > Contexts.

Step 2 Click **Configuration** > **Firewall** > **Service Policy Rules**.

Step 3 You can edit an existing rule or create a new one:

- For an existing rule, choose the rule and click **Edit**.

  The Edit Service Policy Rule dialog box appears.

- For a new rule, choose **Add** > **Add Service Policy Rule**.

  The Add Service Policy Rule Wizard - Service Policy dialog box appears. Complete the Service Policy and Traffic Classification Criteria dialog boxes. See the "Adding a Service Policy Rule for Through Traffic" section on page 23-7 for more information. Click **Next** to show the Add Service Policy Rule Wizard - Rule Actions dialog box.

Step 4 Click the **Intrusion Prevention** tab.

You can also set other feature actions for the same traffic using the other tabs.

Step 5 Check the **Enable IPS for this traffic flow** check box.

Step 6 In the Mode area, click **Inline Mode** or **Promiscuous Mode**.

See the "Operating Modes" section on page 30-2 for more details.

Step 7 In the If IPS Card Fails area, click **Permit traffic** or **Close traffic**.

The Close traffic option sets the adaptive security appliance to block all traffic if the AIP SSM/SSC is unavailable.

The Permit traffic option sets the adaptive security appliance to allow all traffic through, uninspected, if the AIP SSM/SSC is unavailable.

Step 8 (AIP SSM Only) From the IPS Sensor to use drop-down list, choose a virtual sensor name.

If you use virtual sensors on the AIP SSM only, you can specify a sensor name using this option. If you use multiple context mode on the security appliance, you can only specify sensors that you assigned to the context (see the "Assigning Virtual Sensors to Security Contexts (AIP SSM Only)" section on page 30-7). If you do not specify a sensor name, then the traffic uses the default sensor. In multiple context mode, you can specify a default sensor for the context. In single mode or if you do not specify a default sensor in multiple mode, the traffic uses the default sensor that is set on the AIP SSM.

Step 9 Click **OK**.

Resetting the AIP SSM/SSC Password

You can use ASDM to reset the AIP SSM/SSC password to the default if the AIP SSM/SSC 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 the "AIP SSM/SSC Task Overview" section on page 30-6 for information about accessing IDM from ASDM.

Resetting the AIP SSM/SSC password causes the AIP SSM/SSC to reboot. IPS services are not available while the AIP SSM/SSC is rebooting.

To reset the AIP SSM/SSC 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/SSC 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 IPS Version 6.0 or later on the AIP SSM/SSC.

Step 3 Click **Close** to close the dialog box.

Feature History for the AIP SSM/SSC

Table 30-1 lists the release history for this feature.

Table 30-1 Feature History for the AIP SSM/SSC

Feature Name	Releases	Feature Information
AIP SSM	7.0(1)	The AIP SSM was introduced. The following command was introduced: **ips**.
Virtual sensors	8.0(2)	Virtual sensor support was introduced. Virtual sensors let you configure multiple security policies on the AIP SSM. The following command was introduced: **allocate-ips**.
AIP SSC for the ASA 5505	8.2(1)	The AIP SSC was introduced. The following commands were introduced: **allow-ssc-mgmt**, **hw-module module ip**, and **hw-module module allow-ip**.

CHAPTER 31

Configuring Trend Micro Content Security

Note The ASA 5580 does not support the CSC SSM feature.

This chapter describes how to configure the CSC SSM using the CSC Setup Wizard in ASDM and the CSC SSM GUI, and includes the following sections:

Information About the CSC SSM

The ASA 5500 series 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 by scanning the FTP, HTTP, POP3, and SMTP packets that you configure the security appliance to send to it.

For more information about the CSC SSM, see the following URL:

http://www.cisco.com/en/US/products/ps6823/index.html

Licensing Requirements for the CSC SSM

The following table shows the licensing requirements for this feature:

Model	License Requirement
ASA 5505	No support.
ASA 5510	Security Plus License: 2 contexts. *Optional license: 5 contexts.*
ASA 5520	Base License: 2 contexts. *Optional licenses: 5, 10, or 20 contexts.*
ASA 5540	Base License: 2 contexts. *Optional licenses: 5, 10, 20, or 50 contexts.*

For the ASA 5510, 5520, and 5540:

- 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 Security Plus License, the additional features enabled by default are SMTP anti-spam, SMTP content filtering, POP3 anti-spam, URL blocking, and URL filtering.

Prerequisites for the CSC SSM

The CSC SSM has the following prerequisites:

- A CSC SSM card must be installed in the security appliance.
- A Product Authorization Key (PAK) for use in registering the CSC SSM.
- Activation keys that you receive by e-mail after you register the CSC SSM.
- The management port of the CSC SSM must be connected to your network to allow management and automatic updates of the CSC SSM software.
- The CSC SSM management port IP address must be accessible by the hosts used to run ASDM.
- You must obtain the following information to use in configuring the CSC SSM:
  - The CSC SSM management port IP address, netmask, and gateway IP address.
  - DNS server IP address.
  - HTTP proxy server IP address (needed only if your security policies require the use of a proxy server for HTTP access to the Internet).
  - Domain name and hostname for the CSC SSM.
  - An e-mail address and an SMTP server IP address and port number for e-mail notifications.
  - IP addresses of hosts or networks that are allowed to manage the CSC SSM. The IP addresses for the CSC SSM management port and the security appliance management interface can be in different subnets.
  - Password for the CSC SSM.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context modes. In multiple-context mode, all panes under the CSC Setup node are available *only* in the admin context. You can restore the default password only in multiple-context mode in the system context.

Firewall Mode Guidelines

Supported in routed and transparent firewall modes.

Failover Guidelines

Does not support sessions in Stateful Failover. The CSC SSM does not maintain connection information, and therefore cannot provide the failover unit with the required information. The connections that a CSC SSM is scanning are dropped when the security appliance in which the CSC SSM is installed fails. When the standby security appliance becomes active, it forwards the scanned traffic to the CSC SSM and the connections are reset.

IPv6 Guidelines

Does not support IPv6.

Model Guidelines

Supported on the ASA 5510, ASA 5520, and ASA 5540 only.

Default Settings

Table 31-1 lists the default settings for the CSC SSM.

Table 31-1 Default CSC SSM Parameters

Parameter	Default
FTP inspection on the security appliance	Enabled
All features included in the license(s) that you have purchased	Enabled

CSC SSM Setup

The CSC Setup Wizard lets you configure basic operational parameters for the CSC SSM. You must complete this wizard at least once before you can configure options in each screen separately. After you complete the CSC Setup Wizard, you can modify each screen individually without using this wizard again.

Additionally, you cannot access the panes under Configuration > Trend Micro Content Security > CSC Setup or under Monitoring > Trend Micro Content Security > Content Security until you complete the CSC Setup Wizard. If you try to access these panes before completing this wizard, a dialog box appears and lets you access the wizard directly to complete the configuration.

This section includes the following topics:

- Activation/License, page 31-4
- IP Configuration, page 31-5
- Host/Notification Settings, page 31-5
- Management Access Host/Networks, page 31-6
- Password, page 31-7
- Restoring the Default Password, page 31-7
- Wizard Setup, page 31-8

Activation/License

The Activation/License pane lets you review or renew activation codes for the CSC SSM Base License and the 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. Links to the licensing status pane and the CSC UI home pane appear at the bottom of this window. The serial number for the assigned license is filled in automatically.

To review license status or renew a license, perform the following steps:

Step 1 The Activation/License pane shows the following display-only information for the Base License and the Plus License:

- The name of the component.
- The activation code for the corresponding Product field.
- The status of the license. If the license is valid, the expiration date appears. If the expiration date has passed, this field indicates that the license has expired.
- The maximum number of network devices that the Base License supports. The Plus License does not affect the number of network devices supported; therefore, the Nodes field does not appear in the Plus License area.

Step 2 To review license status or renew your license, click the link provided.

Step 3 To go to the CSC home pane in ASDM, click the link provided.

What to Do Next

See the "IP Configuration" section on page 31-5.

IP Configuration

The IP Configuration pane lets you configure management access for the CSC SSM, the DNS servers it should use, and a proxy server for retrieving CSC SSM software updates.

To configure management access and other related details for the CSC SSM, perform the following steps:

Step 1 Set the following parameters for management access to the CSC SSM:

- Enter the IP address for management access to the CSC SSM.
- Enters the netmask for the network containing the management IP address of the CSC SSM.
- Enter the IP address of the gateway device for the network that includes the management IP address of the CSC SSM.

Step 2 Set parameters of the DNS servers for the network that includes the management IP address of the CSC SSM.

- Enter the IP address of the primary DNS server.
- (Optional) Enter the IP address of the secondary DNS server.

Step 3 (Optional) Enter parameters for an 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, leave the fields in this group blank.

- Enter the IP address of the proxy server.
- Enter the listening port of the proxy server.

What to Do Next

See the “Host/Notification Settings” section on page 31-5.

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 to be excluded from detailed scanning.

To configure host and notification settings, perform the following steps:

Step 1 In the Host and Domain Names area, set the hostname and domain name of the CSC SSM.

Step 2 In the Incoming E-mail Domain Name area, set the trusted incoming e-mail domain name for SMTP-based e-mail. The CSC SSM scans SMTP e-mail sent to this domain. The types of threats that the CSC SSM scans for depend on the license that 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 > CSC Setup > Mail,** and then click one of the links. After logging in to the CSC SSM, choose **Mail (SMTP)** > **Configuration**, and then click the **Incoming Mail** tab.

Step 3 Configure the following settings for e-mail notification of events:

- The administrator e-mail address for the account to which notification e-mails should be sent.
- The IP address of the SMTP server.
- The port to which the SMTP server listens.

What to Do Next

See the "Management Access Host/Networks" section on page 31-6.

Management Access Host/Networks

The Management Access Host/Networks pane lets you specify the hosts and networks for which management access to the CSC SSM is permitted. You must specify at least one permitted host or network, up to a maximum of eight permitted hosts or networks.

To specify hosts and networks for which management access to the CSC SSM is allowed, perform the following steps:

Step 1 Enter the IP address of a host or network that you want to add to the Selected Hosts/Network list.

Step 2 Enter the netmask for the host or network that you specified in the IP Address field.

Note 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.

The Selected Hosts/Networks list displays the hosts or networks trusted for management access to the CSC SSM.

Step 3 To add the host or network that you specified in the IP Address field in the Selected Hosts/Networks list, click **Add**.

Step 4 To remove a host or network from the Selected Hosts/Networks list, choose an entry from the list and click **Delete**.

What to Do Next

See the "Password" section on page 31-7.

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; however, 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. As a result, ASDM displays a confirmation dialog box that you must respond to before the password is changed.

Tip Whenever the connection to the CSC SSM is dropped, you can reestablish it. To do so, click the **Connection to Device** icon on the status bar to display the Connection to Device dialog box, and then click **Reconnect**. ASDM prompts you for the CSC SSM password, which is the new password that you have defined.

Passwords must be 5 - 32 characters long.

Passwords appears as asterisks when you type them.

Note The default password is “cisco.”

To change the password required for management access to the CSC SSM, perform the following steps:

Step 1 In the Old Password field, enter the current password for management access to the CSC SSM.

Step 2 In the New Password field, enter the new password for management access to the CSC SSM.

Step 3 In the Confirm New Password field, reenter the new password for management access to the CSC SSM.

What to Do Next

If required, see the “Restoring the Default Password” section on page 31-7.

See the “Wizard Setup” section on page 31-8.

Restoring the Default Password

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). If the CSC password-reset policy has been set to “Denied,” then you cannot reset the password through the ASDM CLI. To change this policy, you must access the CSC SSM through the security appliance CLI by entering the **session** command. For more information, see the *Trend Micro InterScan for Cisco CSC SSM Administrator Guide*.

Note This option does not appear in the menu if an SSM is not installed.

To reset the CSC SSM password to the default value, 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 value.

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 8.0(2) software on the security appliance and the most recent Version 6.1.x software 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.

What to Do Next

See the "Password" section on page 31-7.

Wizard Setup

The Wizard Setup screen lets you start the CSC Setup Wizard. To start the CSC Setup Wizard, click **Launch Setup Wizard**.

Before you can directly access any of the other screens under CSC Setup, you must complete the CSC Setup Wizard. This wizard includes the following screens:

- CSC Setup Wizard Activation Codes Configuration, page 31-8
- CSC Setup Wizard IP Configuration, page 31-9
- CSC Setup Wizard Host Configuration, page 31-9
- CSC Setup Wizard Management Access Configuration, page 31-10
- CSC Setup Wizard Password Configuration, page 31-10
- CSC Setup Wizard Traffic Selection for CSC Scan, page 31-10
- CSC Setup Wizard Summary, page 31-12

After you complete the CSC Setup Wizard once, you can change any settings in screens related to the CSC SSM without using the CSC Setup Wizard again.

CSC Setup Wizard Activation Codes Configuration

To display the activation codes that you have entered to enable features on the CSC SSM, perform the following steps:

Choose **Configuration** > **Trend Micro Content Security** > **CSC Setup** > **Activation/License**.

The activation code settings that you have made appear on this screen, according to the type of license you have, as follows:

- The activation code for the Base License appears. The Base License includes anti-virus, anti-spyware, and file blocking.
- The activation code for the Plus License appears, if you have entered one. If not, this field is blank. The Plus License includes anti-spam, anti-phishing, content filtering, URL blocking and filtering, and web reputation.

What to Do Next

See the "CSC Setup Wizard IP Configuration" section on page 31-9.

CSC Setup Wizard IP Configuration

To display the IP configuration settings that you have entered for the CSC SSM, perform the following steps:

Choose **Configuration** > **Trend Micro Content Security** > **CSC Setup** > **IP Configuration**.

The IP configuration settings that you have entered for the CSC SSM appear, including the following:

- The IP address for the management interface of the CSC SSM.
- The network mask for the management interface of the CSC SSM that you have selected from the drop-down list.
- The IP address of the gateway device for the network that contains the CSC SSM management interface.
- The primary DNS server IP address.
- The secondary DNS server IP address (if configured).
- The proxy server (if configured).
- The proxy port (if configured).

What to Do Next

See the "CSC Setup Wizard IP Configuration" section on page 31-9.

CSC Setup Wizard Host Configuration

To display the host configuration settings that you have entered for the CSC SSM, perform the following steps:

Choose **Configuration** > **Trend Micro Content Security** > **CSC Setup** > **Host Configuration**.

The host configuration settings that you have entered for the CSC SSM appear, including the following:

- The hostname of the CSC SSM.
- The name of the domain in which the CSC SSM resides.
- The domain name for incoming e-mail.
- The e-mail address of the domain administrator.
- The IP address of the e-mail server.
- The port number through which you connect to the CSC SSM.

What to Do Next

See the "CSC Setup Wizard Management Access Configuration" section on page 31-10.

CSC Setup Wizard Management Access Configuration

To display the subnet and host settings that you have entered to grant access to the CSC SSM, perform the following steps:

Step 1 Choose **Configuration** > **Trend Micro Content Security** > **CSC Setup** > **Management Access Configuration**.

The management access configuration settings that you have entered for the CSC SSM appear, including the following:

- The IP address for networks and hosts that are allowed to connect to the CSC SSM.
- The network mask for networks and hosts that are allowed to connect to the CSC SSM that you have selected from the drop-down list.

Step 2 To add the IP address of the networks and hosts that you want to allow to connect to the CSC SSM, click **Add**.

Step 3 To remove the IP address of a network or host whose ability to connect to the CSC SSM you no longer want, click **Delete**.

The Selected Hosts/Networks table lists the IP addresses of networks and hosts whose connection to the CSC SSM you have added.

What to Do Next

See the "CSC Setup Wizard Password Configuration" section on page 31-10.

CSC Setup Wizard Password Configuration

To change the password required for management access to the CSC SSM, perform the following steps:

Step 1 Choose **Configuration** > **Trend Micro Content Security** > **CSC Setup** > **Password**.

Step 2 In the Old Password field, enter the current password for management access to the CSC SSM.

Step 3 In the New Password field, enter the new password for management access to the CSC SSM.

Step 4 In the Confirm New Password field, reenter the new password for management access to the CSC SSM.

What to Do Next

See the "CSC Setup Wizard Traffic Selection for CSC Scan" section on page 31-10.

CSC Setup Wizard Traffic Selection for CSC Scan

To display the settings that you have made to select traffic for CSC scanning, perform the following steps:

Step 1 Choose **Configuration** > **Trend Micro Content Security** > **CSC Setup** > **Traffic Selection for CSC Scan**.

The traffic selection for CSC scanning configuration settings that you have entered for the CSC SSM appear, including the following:

- The interface to the CSC SSM that you have chosen from the drop-down list.
- The source of network traffic for the CSC SSM to scan.
- The destination of network traffic for the CSC SSM to scan.
- The source or destination service for the CSC SSM to scan.

Step 2 Do one of the following:

- To specify additional traffic details for CSC scanning, click **Add**. For more information, see "Specifying Traffic for CSC Scanning" section on page 31-11.
- To modify additional traffic details for CSC scanning, click **Edit**. For more information, see "Specifying Traffic for CSC Scanning" section on page 31-11.
- To remove additional traffic details for CSC scanning, click **Delete**.

Specifying Traffic for CSC Scanning

To define, modify, or remove additional settings for selecting traffic for CSC scanning, perform the following steps:

Step 1 In the Traffic Selection for CSC Scan screen, click **Specify traffic for CSC Scan**.

The Specify traffic for CSC Scan dialog box appears.

Step 2 Choose the type of interface to the CSC SSM from the drop-down list. Available settings are global (all interfaces), inside, management, and outside.

Step 3 Choose the source of network traffic for the CSC SSM to scan from the drop-down list.

Step 4 Choose the destination of network traffic for the CSC SSM to scan from the drop-down list.

Step 5 Choose the type of service for the CSC SSM to scan from the drop-down list.

Step 6 Enter a description for the network traffic that you define for the CSC SSM to scan.

Step 7 Specify whether or not to allow the CSC SSM to scan network traffic if the CSC card fails. Choose one of the following options:

- To allow traffic through without being scanned, click **Permit**.
- To prevent traffic from going through without being scanned, click **Close**.

Step 8 Click **OK** to save your settings.

The added traffic details appear on the CSC Setup Wizard Traffic selection for CSC Scan screen.

Step 9 Click **Cancel** to discard these settings and return to the CSC Setup Wizard Traffic selection for CSC Scan screen. If you click **Cancel**, ASDM displays a dialog box to confirm your decision.

What to Do Next

See the "CSC Setup Wizard Summary" section on page 31-12.

CSC Setup Wizard Summary

To review the settings that you have made with the CSC Setup Wizard, perform the following steps:

Step 1 Choose **Configuration** > **Trend Micro Content Security** > **CSC Setup** > **Summary**.

The CSC Setup Wizard Summary screen shows the following display-only settings:

- The settings that you made in the Activation Codes Configuration screen, including the Base License activation code and the Plus License activation code, if you entered one. If not, this field is blank.
- The settings that you made in the IP Configuration screen, including the following information:
  - IP address and netmask for the management interface of the CSC SSM.
  - IP address of the gateway device for the network that includes the CSC SSM management interface.
  - Primary DNS server IP address.
  - Secondary DNS server IP address (if configured).
  - Proxy server and port (if configured).
- The settings that you made in the Host Configuration screen, including the following information:
  - Hostname of the CSC SSM.
  - Domain name for the domain that includes the CSC SSM.
  - Domain name for incoming e-mail.
  - Administrator e-mail address.
  - E-mail server IP address and port number.
- The settings that you made on the Management Access Configuration screen. The drop-down list includes the hosts and networks from which the CSC SSM allows management connections.
- Indicates whether or not you have changed the password in the Password Configuration screen.

Step 2 (Optional) Click **Back** to return to the previous screens of the CSC Setup Wizard to change any settings.

Note The Next button is dimmed; however, if you click **Back** to access any of the preceding screens in this wizard, click **Next** to return to the Summary screen.

Step 3 Click **Finish** to complete the CSC Setup Wizard and save all settings that you have specified. After you click **Finish**, you can change any settings related to the CSC SSM without using the CSC Setup Wizard again.

A summary of the status of commands that were sent to the device appears.

Step 4 Click **Close** to close this screen, and then click **Next**.

A message appears indicating that the CSC SSM has been activated and is ready for use.

Step 5 (Optional) Click **Cancel** to exit the CSC Setup Wizard without saving any of the selected settings. If you click **Cancel**, a dialog box appears to confirm your decision.

What to Do Next

See the "Using the CSC SSM GUI" section on page 31-13.

Using the CSC SSM GUI

This section describes how to configure features using the CSC SSM GUI, and includes the following topics:

- Web, page 31-13
- Mail, page 31-14
- File Transfer, page 31-15
- Updates, page 31-16

Web

Note To access the CSC SSM, you must reenter the CSC SSM password. Sessions in the CSC SSM browser 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 again, because one session is already open.

To view whether or not web-related features are enabled and access the CSC SSM GUI for configuring these features, perform the following steps:

Step 1 Choose **Configuration** > **Trend Micro Content Security** > **Web**.

The URL Blocking and Filtering area is display-only and shows whether or not URL blocking is enabled on the CSC SSM.

Step 2 Click **Configure URL Blocking** to open a screen for configuring URL blocking on the CSC SSM.

The URL Filtering area is display-only and shows whether or not URL filtering is enabled on the CSC SSM.

Step 3 Click **Configure URL Filtering** to open a screen for configuring URL filtering rules on the CSC SSM.

The File Blocking area is display-only and shows whether or not URL file blocking is enabled on the CSC SSM.

Step 4 Click **Configure File Blocking** to open a screen for configuring file blocking settings on the CSC SSM.

The HTTP Scanning area is display-only and shows whether or not HTTP scanning is enabled on the CSC SSM.

Step 5 Click **Configure Web Scanning** to open a screen for configuring HTTP scanning settings on the CSC SSM.

The Web Reputation area is display-only and shows whether or not the Web Reputation service is enabled on the CSC SSM.

Step 6 Click **Configure Web Reputation** to open a screen for configuring the Web Reputation service on the CSC SSM.

What to Do Next

See the “Mail” section on page 31-14.

Mail

The Mail pane lets you see whether or not e-mail-related features are enabled and lets you access the CSC SSM GUI to configure these features. To configure e-mail related features, choose **Configuration > Trend Micro Content Security > Mail**.

This section includes the following topics:

- SMTP Tab, page 31-14
- POP3 Tab, page 31-15

SMTP Tab

Note To access the CSC SSM, you must reenter the CSC SSM password. Sessions in the CSC SSM browser 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 again, because one session is already open.

To configure SMTP scanning, perform the following steps:

Step 1 Click the **SMTP** Tab.

Step 2 The Incoming Scan area is display-only and shows whether or not the incoming SMTP scanning feature is enabled on the CSC SSM. Click **Configure Incoming Scan** to open a screen for configuring incoming SMTP scan settings on the CSC SSM.

Step 3 The Outgoing Scan area is display-only and shows whether or not the outgoing SMTP scanning feature is enabled on the CSC SSM. Click **Configure Outgoing Scan** to open a screen for configuring outgoing SMTP scan settings on the CSC SSM.

Step 4 The Incoming Filtering area is display-only and shows whether or not content filtering for incoming SMTP e-mail is enabled on the CSC SSM. Click **Configure Incoming Filtering** to open a screen for configuring incoming SMTP e-mail content filtering settings on the CSC SSM.

Step 5 The Outgoing Filtering area is display-only and shows whether or not content filtering for outgoing SMTP e-mail is enabled on the CSC SSM. Click **Configure Outgoing Filtering** to open a screen for configuring outgoing SMTP e-mail content filtering settings on the CSC SSM.

Step 6 The Anti-spam area is display-only and shows whether or not the SMTP anti-spam feature is enabled on the CSC SSM. Click **Configure Anti-spam** to open a screen for configuring SMTP anti-spam settings, including E-mail Reputation, on the CSC SSM.

POP3 Tab

Note To access the CSC SSM, you must reenter the CSC SSM password. Sessions in the CSC SSM browser 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 again, because one session is already open.

To configure POP3 scanning, perform the following steps:

Step 1 Click the **POP3** Tab.

Step 2 The Scanning area is display-only and shows whether or not POP3 e-mail scanning is enabled on the CSC SSM. Click **Configure Scanning** to open a window for configuring POP3 e-mail scanning on the CSC SSM.

Step 3 The Anti-spam area is display-only and shows whether or not the POP3 anti-spam feature is enabled on the CSC SSM. Click **Configure Anti-spam** to open a window for configuring the POP3 anti-spam feature on the CSC SSM.

Step 4 The Content Filtering area is display-only and shows whether or not POP3 e-mail content filtering is enabled on the CSC SSM. Click **Configure Content Filtering** to open a window for configuring POP3 e-mail content filtering on the CSC SSM.

What to Do Next

See the "File Transfer" section on page 31-15.

File Transfer

The File Transfer pane lets you view whether or not FTP-related features are enabled and lets you access the CSC SSM for configuring FTP-related features.

Note To access the CSC SSM, you must reenter the CSC SSM password. Sessions in the CSC SSM browser 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 again, because one session is already open.

To view the status or configure FTP-related features, perform the following steps:

Step 1 Click the **File Transfer** tab.

The File Scanning area is display-only and shows whether or not FTP file scanning is enabled on the CSC SSM.

Step 2 Click **Configure File Scanning** to open a window for configuring FTP file scanning settings on the CSC SSM.

The File Blocking area is display-only and shows whether or not FTP blocking is enabled on the CSC SSM.

Step 3 Click **Configure File Blocking** to open a window for configuring FTP file blocking settings on the CSC SSM.

What to Do Next

See the "Updates" section on page 31-16.

Updates

The Updates pane lets you view whether or not scheduled updates are enabled and lets you access the CSC SSM for configuring scheduled updates.

Note To access the CSC SSM, you must reenter the CSC SSM password. Sessions in the CSC SSM browser 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 again, because one session is already open.

To view the status or configure scheduled update settings, perform the following steps:

Step 1 Click the **Updates** tab.

The Scheduled Updates area is display-only and shows whether or not scheduled updates are enabled on the CSC SSM.

The Scheduled Update Frequency area displays information about when updates are scheduled to occur, such as "Hourly at 10 minutes past the hour."

The Component area displays names of parts of the CSC SSM software that can be updated.

In the Components area, the Scheduled Updates area is display-only and shows whether or not scheduled updates are enabled for the corresponding components.

Step 2 Click **Configure Updates** to open a window for configuring scheduled update settings on the CSC SSM.

Note If you restart the security appliance, the SSM is not automatically restarted. For more information, see the "Managing SSMs and SSCs" section in the *Cisco ASA 5500 Series Configuration Guide using the CLI*.

Where to Go Next

See the "Monitoring Trend Micro Content Security" section on page 48-1.

Additional References

For additional information related to implementing the CSC SSM, see the following documents:

Related Topic	Document Title
Instructions on use of the CSC SSM GUI. Additional licensing requirements of specific windows available in the CSC SSM GUI. Reviewing the default content security policies in the CSC SSM GUI before modifying them or entering advanced configuration settings.	*Trend Micro InterScan for Cisco CSC SSM Administrator Guide*
Accessing ASDM for the first time and assistance with the Startup Wizard.	*Cisco ASA 5500 Series Adaptive Security Appliance Getting Started Guide*
Assistance with SSM hardware installation and connection to the security appliance.	*Cisco ASA 5500 Series Hardware Installation Guide*
Technical Documentation, Marketing, and Support-related information	See http://www.cisco.com/en/US/products/ps6823/index.html.

Feature History for the CSC SSM

Table 31-2 lists the release history for this feature.

Table 31-2 Feature History for the CSC SSM

Feature Name	Releases	Feature Information
CSC SSM	ASA 7.0(1), ASDM 5.0(1)	The CSC SSM runs Content Security and Control software, which provides protection against viruses, spyware, spam, and other unwanted traffic. The following commands were introduced: **csc {fail-close | fail-open}, hw-module module 1** [**recover** | **reload** | **reset** | **shutdown**], **session, show module** [**all** | *slot* [**details** | **recover**]].
Password reset	ASA 7.2(2), ASDM 5.2(2)	The **hw-module module password-reset** command was introduced.
CSC SSM	ASA 8.1(1), ASDM 6.1(1) and ASA 8.1(2), ASDM 6.1(2)	This feature is not supported.
Syslog format	ASA 8.3(1), ASDM 6.3(1)	CSC syslog format is consistent with the security appliance syslog format. Syslog message explanations have been added to the *Trend Micro InterScan for Cisco CSC SSM Administrator Guide*. The source and destination IP information has been added to the ASDM Log Viewer GUI. All syslog messages include predefined syslog priorities and cannot be configured through the CSC SSM GUI.

PART 4

Configuring VPN

CHAPTER 32

SSL VPN Wizard

SSL VPN Feature

Clientless, browser-based SSL VPN lets users establish a secure, remote-access VPN tunnel to the security appliance using a web browser. After authentication, users access a portal page and can access specific, supported internal resources. The network administrator provides access to resources by users on a group basis. Users have no direct access to resources on the internal network.

The Cisco AnyConnect VPN client provides secure SSL connections to the security appliance for remote users with full VPN tunneling to corporate resouces. Without a previously-installed client, remote users enter the IP address in their browser of an interface configured to accept clientless SSL VPN connections. The security appliance downloads the client that matches the operating system of the remote computer. After downloading, the client installs and configures itself, establishes a secure SSL connection and either remains or uninstalls itself (depending on the security appliance configuration) when the connection terminates. In the case of a previously installed client, when the user authenticates, the security appliance examines the revision of the client, and upgrades the client as necessary.

Fields

- **Clientless SSL VPN Access**—Enables clientless, browser-based connections for specific, supported internal resources through a portal page.
- **Cisco SSL VPN Client (AnyConnect VPN Client)**—Enables SSL VPN client connections for full network access. Enables the security appliance to download the AnyConnect client to remote users.

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 Interface

Provide a Connection name (previously called *tunnel group*), enable an interface for SSL VPN connections, and provide digital certificate information in this window.

Fields

- Connection Name—Provide a connection name for this group of connection-oriented attributes.
- SSL VPN Interface—Specify the interface to allow SSL VPN connections.
- Digital Certificate—Specify a certificate, if any, that the security appliance sends to the remote PC.
  - Certificate—Specify the name of the certificate.
- Connection Group Settings—You can enable the security appliance to display a group alias for this connection on the login page.
  - Connection Group Alias—Specify an alias name for the connection.
  - Display Group Alias list at the login page—Enable to display the group alias.
- Information—Displays information remote users need for establishing SSL VPN connections and ASDM connections.

Modes

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

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

User Authentication

Specify authentication information on this screen.

Fields

- Authenticate using a AAA server group—Enable to let the security appliance contact a remote AAA server group to authenticate the user.
- AAA Server Group Name—Select a AAA server group from the list of pre-configured groups, or click **New** to create a new group.
- Authenticate using the local user database—Add new users to the local database stored on the security appliance.
  - Username—Create a username for the user.
  - Password—Create a password for the user.
  - Confirm Password—Re-type the same password to confirm.
  - Add/Delete—Add or delete the user from the local database.

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

Group policies configure common attributes for groups of users. Create a new group policy or select an existing one to modify.

Fields

- Create new group policy—Enable to create a new group policy. Provide a name for the new policy.
- Modify existing group policy—Select an existing group policy to modify.

Modes

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

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

Bookmark List

Bookmark lists appear on the portal page for Clientless, browser-based connections. SSL VPN client users do not see these bookmarks. Create a new bookmark list on this window.

Fields

- Bookmark List—Select an existing list or click **Manage** to create a new list, or import or export bookmark lists.

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 Pools and Client Image

Provide a range of IP addresses to remote SSL VPN users and identify SSL VPN client images to the security appliance in this window.

Fields

- IP Address Pool—SSL VPN clients receive new IP addresses when they connect to the security appliance. Clientless connections do not require new IP addresses. Address Pools define a range of addresses that remote clients can receive.
- IP Address Pool—Select an existing IP Address Pool, or click **New** to create a new pool.
- AnyConnect VPN Client Image Location—Identify to the security appliance files in flash memory that are SSL VPN client images. Click Browse to locate images on your local PC.
  - Location—Provide the path and filename of a valid SSL VPN client image located in flash memory.
  - Download Latest AnyConnect VPN Client form CCO—Click this link to go to the Software Download page for the latest client image.

Modes

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

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

Summary

Provides a summary of your selections from the previous wizard windows.

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

VPN

This chapter describes how to configure a basic VPN connection using the IPsec VPN wizard. It includes the following sections:

- VPN Wizard
- VPN Tunnel Type
- Remote Site Peer
- IKE Policy
- IPsec Rule
- Hosts and Networks
- Remote Access Client
- VPN Client Authentication Method and Name
- Client Authentication
- New Authentication Server Group
- User Accounts
- Address Pool
- Attributes Pushed to Client
- IPsec Settings (Optional)
- Summary

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

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 panes 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

Use the VPN Tunnel Type pane 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 panes 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 panes that let you enter the attributes a remote access VPN requires.
- VPN Tunnel Interface—Choose 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

Use the Remote Site Peer pane for the following tasks:

1. Providing the IP address of the remote IPsec peer that terminates this VPN tunnel.
2. Selecting and configuring an authentication method.
3. Creating a connection policy (tunnel group).

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 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—Displays the algorithm for signing digital certificates, rsa-sig for RSA.
  - Certificate Name—Choose 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 connection policy can specify authentication, authorization, and accounting servers, a default group policy, and IKE attributes. A policy that you configure with this VPN wizard specifies an authentication method, and uses the security appliance Default Group Policy.

  By default, ASDM populates this field 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

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 pane 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—Choose 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.
- Diffie-Hellman Group—Choose 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).

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 Rule

Use this IPsec Rule pane 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—Choose 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—Choose 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.

- Enable Perfect Forwarding Secrecy (PFS)—Specify whether to use Perfect Forward Secrecy, and the size of the numbers to use, in generating Phase 2 IPsec keys. 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 PFS is enabled. PFS uses Diffie-Hellman techniques to generate the keys.

  PFS ensures that a session key derived from a set of long-term public and private keys is not compromised if one of the private keys is compromised in the future.

  PFS must be enabled on both sides of the connection.

- Diffie-Hellman 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).

Modes

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

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

Hosts and Networks

Use the Hosts and Networks pane to identify local and remote hosts and networks that can use this LAN-to-LAN IPsec tunnel to send and receive data.

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 pane 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

- Action—Decide whether or not to protect data travelling between the local and remote network.
- Local networks—Select the local hosts and networks.
- Remote networks—Select the remote hosts and networks.
- 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
•	—	•	—	—

Remote Access Client

Use the Remote Access Client pane to identify the type of remote access users this connection serves.

Fields

- Cisco VPN Client Version 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-Proxy—Enables EAP which permits the security appliance to proxy the PPP authentication process to an external RADIUS authentication server.
- Client will send the 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 connection policies. Since each connection policy has its own AAA server group and IP address pools, users can authenticate through methods specific to their policy.

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 Name

Use the VPN Client Authentication Method and Name pane to configure an authentication method and create a connection policy (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 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 Name—Choose the name that identifies the certificate the security appliance sends to the remote peer.
- Certificate Signing Algorithm—Displays the algorithm for signing digital certificates, rsa-sig for RSA.
- 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.

- Name—Type a name to create the record that contains tunnel connection policies for this IPsec connection. A connection policy can specify authentication, authorization, and accounting servers, a default group policy, and IKE attributes. A connection policy 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

Use the Client Authentication pane 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 pane 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 Name—Choose 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

User the New Authentication Server Group pane 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—Choose 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

Use the User Accounts pane 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

Use the Address Pool pane to configure a pool of local IP addresses that the security appliance assigns to remote VPN clients.

Fields

- Name—Displays the name of the connection policy to which the address pool applies. You set this name in the VPN Client Name and Authentication Method pane.
- Pool Name—Select a descriptive identifier for the address pool.
- New...—Click to configure a new address pool.
- 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) Choose 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

Use the Attributes Pushed to Client (Optional) pane 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 connection policy to which the address pool applies. You set this name in the VPN Client Name and Authentication Method pane.
- 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
•	—	•	—	—

IPsec Settings (Optional)

Use the IPsec Settings (Optional) pane 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 pane. 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.
  - 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.
- 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.
- Enable Perfect Forwarding Secrecy (PFS)—Specify whether to use Perfect Forward Secrecy, and the size of the numbers to use, in generating Phase 2 IPsec keys. 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 PFS is enabled. PFS uses Diffie-Hellman techniques to generate the keys.

  PFS ensures that a session key derived from a set of long-term public and private keys is not compromised if one of the private keys is compromised in the future.

  PFS must be enabled on both sides of the connection.
  - Diffie-Hellman 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).

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 pane 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 pane.

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
•	—	•	—	—

CHAPTER 34

Configuring Certificates

This chapter describes how to configure digital certificates, and includes the following sections:

Information About Digital Certificates

Digital certificates provide digital identification for authentication. A digital certificate includes information that identifies a device or user, such as the name, serial number, company, department, or IP address. CAs are trusted authorities that "sign" certificates to verify their authenticity, thereby guaranteeing the identity of the device or user. CAs issue digital certificates in the context of a PKI, which uses public-key or private-key encryption to ensure security.

For authentication using digital certificates, at least one identity certificate and its issuing CA certificate must exist on an security appliance. This configuration allows multiple identities, roots, and certificate hierarchies. Descriptions of several different types of available digital certificates follow:

- A *CA certificate* is used to sign other certificates. It is self-signed and called a *root certificate*. A certificate that is issued by another CA certificate is called a *subordinate certificate*. For more information, see the "Configuring CA Certificate Authentication" section on page 34-2.
- CAs also issue *identity certificates*, which are certificates for specific systems or hosts. For more information, see the "Configuring Identity Certificates Authentication" section on page 34-8.
- *Code-signer certificates* are special certificates that are used to create digital signatures to sign code, with the signed code itself revealing the certificate origin. For more information, see the "Configuring Code Signer Certificates" section on page 34-14.

The local CA integrates an independent certificate authority feature on the security appliance, deploys certificates, and provides secure revocation checking of issued certificates. The local CA provides a secure, configurable, in-house authority for certificate authentication with user enrollment through a website login page. For more information, see the "Authenticating Using the Local CA" section on page 34-16, the "Managing User Certificates" section on page 34-22, and the "Managing the User Database" section on page 34-19.

Note CA certificates and identity certificates apply to both site-to-site VPN connections and remote access VPN connections. Procedures in this document refer to remote access VPN use in the ASDM GUI.

Licensing Requirements for Digital Certificates

The following table shows the licensing requirements for this feature:

Model	License Requirement
All models	Base License.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

Failover Guidelines

Does not support replicating sessions in Stateful Failover.

IPv6 Guidelines

Supports IPv6.

Configuring CA Certificate Authentication

The CA Certificates pane displays the available certificates, identified by the issued to and issued by CA server, the date that the certificate expires, the associated trustpoints, and the certificate usage or purpose. In the CA Certificates pane, you can perform the following tasks:

- Authenticate self-signed or subordinate CA certificates.
- Install CA certificates on the security appliance.
- Create a new certificate configuration.
- Edit an existing certificate configuration.

- Obtain a CA certificate manually and import it.
- Have the security appliance use SCEP to contact the CA, and then automatically obtain and install the certificate.
- Display details and issuer information for a selected certificate.
- Access the CRL for an existing CA certificate.
- Remove the configuration of an existing CA certificate.
- Save the new or modified CA certificate configuration.
- Discard any changes and return the certificate configuration to the original settings.

This section includes the following topics:

Adding or Installing a CA Certificate

You can add a new certificate configuration from an existing file, by manually pasting a certificate in PEM format, or by automatic enrollment using SCEP. SCEP is a secure messaging protocol that requires minimal user intervention and lets you enroll and install certificates using only the VPN Concentrator Manager.

To add or install a CA certificate, perform the following steps:

Step 1 In the main ASDM application window, choose **Configuration** > **Remote Access VPN** > **Certificate Management** > **CA Certificates**.

Step 2 Click **Add**.

The Install Certificate dialog box appears. The selected trustpoint name appears in read-only format.

Step 3 To add a certificate configuration from an existing file, click the **Install from a file** radio button (this is the default setting).

Step 4 Enter the path and file name, or click **Browse** to search for the file. Then click **Install Certificate**.

Step 5 To enroll manually, click the **Paste certificate in PEM format** radio button.

Step 6 Copy and paste the PEM format (base64 or hexadecimal) certificate into the area provided, then click **Install Certificate**.

Step 7 To enroll automatically, click the **Use SCEP** radio button. The security appliance contacts the CA using SCEP, obtains the certificates, and installs them on the device. To use SCEP, you must enroll with a CA that supports SCEP, and you must enroll via the Internet. Automatic enrollment using SCEP requires that you provide the following information:

- The path and file name of the certificate to be automatically installed.
- The maximum number of minutes to retry certificate installation.The default is one minute.
- The number of retries for installing a certificate. The default is zero, which indicates unlimited retries within the retry period.

Step 8 To display additional configuration options for new and existing certificates, click **More Options**.

The Configuration Options for CA Certificates pane appears.

Step 9 To continue, see the "Configuring CA Certificates for Revocation" section on page 34-5.

Editing or Removing a CA Certificate Configuration

To change or remove an existing CA certificate configuration, perform the following steps:

Step 1 To change an existing CA certificate configuration, select it, and then click **Edit**.

The Edit Options for CA Certificates pane appears. To change any of these settings, see the following sections for procedures:

- "Configuring CA Certificates for Revocation" section on page 34-5
- "Configuring CRL Retrieval Policy" section on page 34-5
- "Configuring CRL Retrieval Methods" section on page 34-6
- "Configuring OCSP Rules" section on page 34-6
- "Configuring Advanced CRL and OCSP Settings" section on page 34-7

Step 2 To remove a CA certificate configuration, select it, and then click **Delete**.

Note After you delete a certificate configuration, it cannot be restored. To recreate the deleted certificate, click **Add** to reenter all of the certificate configuration information.

Showing CA Certificate Details

To show detailed information about the selected CA certificate, click **Show Details** to display the Certificate Details dialog box, which includes the following three *display-only* tabs:

- The General tab displays the values for type, serial number, status, usage, public key type, CRL distribution point, the times within which the certificate is valid, and associated trustpoints. The values apply to both available and pending status.
- The Issued to tab displays the X.500 fields of the subject DN or certificate owner and their values. The values apply only to available status.
- The Issued by tab displays the X.500 fields of the entity granting the certificate. The values apply only to available status.

Requesting a CRL

To update the current version of the CRL, click **Request CRL**. CRL updates provide the current status of certificate users. If the request fails, an error message appears. The CRL is updated and regenerated automatically until it expires; clicking **Request CRL** forces an immediate CRL file update and regeneration.

Configuring CA Certificates for Revocation

To configure CA certificates for revocation, perform the following steps:

Step 1 In the Configuration Options for CA Certificates pane, click the **Revocation Check** tab.

Step 2 To disable revocation checking of certificates, click the **Do not check certificates for revocation** radio button.

Step 3 To select one or more revocation checking methods (CRL or OCSP), click the **Check certificates for revocation** radio button.

Step 4 In the Revocation Methods area, available methods appear on the left. Click **Add** to move a method to the right and make it available. Click **Move Up** or **Move Down** to change the method order.

The methods you choose are implemented in the order in which you add them. If a method returns an error, the next revocation checking method activates.

Step 5 Check the **Consider certificate valid if revocation checking returns errors** check box to ignore revocation checking errors during certificate validation.

Step 6 Click **OK** to close the Revocation Check tab. Alternatively, to continue, see the "Configuring CRL Retrieval Policy" section on page 34-5.

Configuring CRL Retrieval Policy

To configure the CRL retrieval policy, perform the following steps:

Step 1 In the Configuration Options for CA Certificates pane, click the **CRL Retrieval Policy** tab.

Step 2 Check the **Use CRL Distribution Point from the certificate** check box to direct revocation checking to the CRL distribution point from the certificate being checked.

Step 3 Check the **Use Static URLs configured below** check box to list specific URLs to be used for CRL retrieval. The URLs you select are implemented in the order in which you add them. If an error occurs with the specified URL, the next URL in order is taken.

Step 4 In the Static Configuration area, click **Add**.

The Add Static URL dialog box appears.

Step 5 In the URL field, enter the static URL to use for distributing the CRLs, and then click **OK**.

The URL that you entered appears in the Static URLs list.

Step 6 To change the static URL, select it, and then click **Edit**.

Step 7 To remove an existing static URL, select it, and then click **Delete**.

Step 8 To change the order in which the static URLs appear, click **Move Up** or **Move Down**.

Step 9 Click **OK** to close this tab. Alternatively, to continue, see the "Configuring CRL Retrieval Methods" section on page 34-6.

Configuring CRL Retrieval Methods

To configure CRL retrieval methods, perform the following steps:

Step 1 In the Configuration Options for CA Certificates pane, click the **CRL Retrieval Methods** tab.

Step 2 Choose one of the following three retrieval methods:

- To enable LDAP for CRL retrieval, check the **Enable Lightweight Directory Access Protocol (LDAP)** check box. With LDAP, CRL retrieval starts an LDAP session by connecting to a named LDAP server, accessed by a password. The connection is on TCP port 389 by default. Enter the following required parameters:
  - Name
  - Password
  - Confirm Password
  - Default Server (server name)
  - Default Port (389)
- To enable HTTP for CRL retrieval, check the **Enable HTTP** check box.
- To enable SCEP for CRL retrieval, check the **Enable Simple Certificate Enrollment Protocol (SCEP)** check box.

Step 3 Click **OK** to close this tab. Alternatively, to continue, see the "Configuring OCSP Rules" section on page 34-6.

Configuring OCSP Rules

The security appliance examines OCSP rules in priority order, and applies the first one that matches. X.509 digital certificates are an alternative to using CRLs.

Note Make sure that you have configured a certificate map before you try to add OCSP rules. If a certificate map has not been configured, an error message appears. To configure a certificate map, choose **Configuration** > **Network (Client) Access, Advanced** > **IPSec** > **Certificate to Connection Profile Maps** > **Rules** > **Add**.

To configure OCSP rules for obtaining revocation status of an X.509 digital certificate, perform the following steps:

Step 1 In the Configuration Options for CA Certificates pane, click the **OCSP Rules** tab.

Step 2 Choose the certificate map to match to this OCSP rule. Certificate maps match user permissions to specific fields in a certificate. The name of the CA that the security appliance uses to validate responder certificates appears in the Certificate field. The priority number for the rule appears in the Index field. The URL of the OCSP server for this certificate appears in the URL field.

Step 3 To add a new OCSP rule, click **Add**.

The Add OCSP Rule dialog box appears.

Step 4 Choose the certificate map to use from the drop-down list.

Step 5 Choose the certificate to use from the drop-down list.

Step 6 Enter the priority number for the rule.

Step 7 Enter the URL of the OCSP server for this certificate.

Step 8 When you are done, click **OK** to close this dialog box.

The newly added OCSP rule appears in the list.

Step 9 To edit an existing OCSP rule, select it, and then click **Edit**.

Step 10 To delete an OCSP rule, select it, and then click **Delete**.

Step 11 Click **OK** to close this tab. Alternatively, to continue, see the "Configuring Advanced CRL and OCSP Settings" section on page 34-7.

Configuring Advanced CRL and OCSP Settings

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, because of 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.

To configure additional CRL and OCSP settings, perform the following steps:

Step 1 In the Configuration Options for CA Certificates pane, click the **Advanced** tab.

Step 2 In the CRL Options area, enter the number of minutes between cache refreshes. The default is 60 minutes. The range is 1-1440 minutes. 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.

Step 3 Check the **Enforce next CRL update** check box to require valid CRLs to have a Next Update value that has not expired. Uncheck the **Enforce next CRL update** check box to let valid CRLs with no Next Update value or a Next Update value that has expired.

Step 4 In the OCSP Options area, enter the URL for the OCSP server. The security appliance uses OCSP servers according to the following order:

1. OCSP URL in a match certificate override rule
2. OCSP URL configured in the selected OCSP Options attribute
3. AIA field of a remote user certificate

Step 5 By default, the **Disable nonce extension** check box is checked, which cryptographically binds requests with responses to avoid replay attacks. This process works by matching the extension in the request to that in the response, ensuring that they are the same. Uncheck the **Disable nonce extension** check box if the OCSP server you are using sends pregenerated responses that do not include this matching nonce extension.

Step 6 In the Validation Policy area, choose one of the following options:

- Click the **SSL** radio button or the **IPSec** radio button to restrict the type of remote session that this CA can be used to validate.
- Click the **SSL and IPSec** radio button to let the CA validate both types of sessions.

Step 7 In the Other Options area, choose one of the following options:

- Check the **Accept certificates issued by this CA** check box to indicate that the security appliance should accept certificates from the specified CA.
- Check the **Accept certificates issued by the subordinate CAs of this CA** check box to indicate that the security appliance should accept certificates from the subordinate CA.

Step 8 Click **OK** to close this tab, and then click **Apply** to save your configuration changes.

What to Do Next

See the "Configuring Identity Certificates Authentication" section on page 34-8.

Configuring Identity Certificates Authentication

An identity certificate can be used to authenticate VPN access through the security appliance. In the Identity Certificates Authentication pane, you can perform the following tasks:

- Add or import a new identity certificate.
- Display details of an identity certificate.
- Delete an existing identity certificate.
- Export an existing identity certificate.
- Install an existing identity certificate.
- Enroll for an identity certificate with Entrust.

This section includes the following topics:

- Adding or Importing an Identity Certificate, page 34-9
- Showing Identity Certificate Details, page 34-11
- Deleting an Identity Certificate, page 34-11
- Exporting an Identity Certificate, page 34-11
- Generating a Certificate Signing Request, page 34-12
- Installing Identity Certificates, page 34-13

Adding or Importing an Identity Certificate

To add or import a new identity certificate configuration, perform the following steps:

Step 1 In the main ASDM application window, choose **Configuration > Remote Access VPN > Certificate Management > Identity Certificates**.

Step 2 Click **Add**.

The Add Identity Certificate dialog box appears, with the selected trustpoint name displayed at the top.

Step 3 To import an identity certificate from an existing file, click the **Import the identity certificate from a file** radio button.

Step 4 Enter the passphrase used to decrypt the PKCS12 file.

Step 5 Enter the path name of the file, or click **Browse** to display the Import ID Certificate File dialog box. Find the certificate file, and then click **Import ID Certificate File**.

Step 6 To add a new identity certificate, click the **Add a new identity certificate** radio button.

Step 7 Click **New** to display the Add Key Pair dialog box.

Step 8 To use the default key pair name, click the **Use default keypair name** radio button.

Step 9 To use a new key pair name, click the **Enter a new key pair name** radio button, and type the new name. The security appliance supports multiple key pairs.

Step 10 Choose the modulus size from the drop-down list.

Step 11 Choose the key pair usage by clicking the **General purpose** radio button (default) or **Special** radio button. When you choose the **Special** radio button, the security appliance generates two key pairs, one for signature use and one for encryption use. This selection indicates that two certificates are required for the corresponding identity.

Step 12 Click **Generate Now** to create new key pairs, and then click **Show** to display the Key Pair Details dialog box, which includes the following *display-only* information:

- The name of the key pair whose public key is to be certified.
- The time of day and the date when the key pair is generated.
- The usage of an RSA key pair.
- The modulus size (bits) of the key pairs: 512, 768, 1024, and 2048. The default is 1024.
- The key data, which includes the specific key data in text format.

Step 13 Click **OK** when you are done to close the Key Pair Details dialog box.

Step 14 Choose a certificate subject DN to form the DN in the identity certificate. and then click **Select** to display the Certificate Subject DN dialog box.

Step 15 Choose one or more DN attributes that you want to add from the drop-down list, enter a value, and then click **Add**. Available X.500 attributes for the Certificate Subject DN are the following:

- Common Name (CN)
- Department (OU)
- Company Name (O)
- Country (C)
- State/Province (ST)
- Location (L)

- E-mail Address (EA)

Step 16 Click **OK** when you are done to close the Certificate Subject DN dialog box.

Step 17 To create self-signed certificates, check the **Generate self-signed certificate** check box.

Step 18 To have the identity certificate act as the local CA, check the **Act as local certificate authority and issue dynamic certificates to TLS proxy** check box.

Step 19 To establish additional identity certificate settings, click **Advanced**.

The Advanced Options dialog box appears, with the following three tabs: Certificate Parameters, Enrollment Mode, and SCEP Challenge Password.

Note Enrollment mode settings and the SCEP challenge password are not available for self-signed certificates.

Step 20 Click the **Certificate Parameters** tab, and then enter the following information:

- The FQDN, an unambiguous domain name, to indicate the position of the node in the DNS tree hierarchy.
- The e-mail address associated with the identity certificate.
- The security appliance IP address on the network in four-part, dotted-decimal notation.
- To add the security appliance serial number to the certificate parameters, check the **Include serial number of the device** check box.

Step 21 Click the **Enrollment Mode** tab, and then enter the following information:

- Choose the enrollment method by clicking the **Request by manual enrollment** radio button or the **Request from a CA** radio button.
- The enrollment URL of the certificate to be automatically installed through SCEP.
- The maximum number of minutes allowed to retry installing an identity certificate. The default is one minute.
- The maximum number of retries allowed for installing an identity certificate. The default is zero, which indicates an unlimited number of retries within the retry period.

Step 22 Click the **SCEP Challenge Password** tab, and then enter the following information:

- The SCEP password
- The SCEP password confirmation

Step 23 Click **OK** when you are done to close the Advanced Options dialog box.

Step 24 Click **Add Certificate** in the Add Identity Certificate pane.

The new identity certificate appears in the Identity Certificates list.

Step 25 Click **Apply** to save the new identity certificate configuration.

Showing Identity Certificate Details

To show detailed information about the selected identity certificate, click **Show Details** to display the Certificate Details dialog box, which includes the following three *display-only* tabs:

- The General tab displays the values for type, serial number, status, usage, public key type, CRL distribution point, the times within which the certificate is valid, and associated trustpoints. The values apply to both available and pending status.
- The Issued to tab displays the X.500 fields of the subject DN or certificate owner and their values. The values apply only to available status.
- The Issued by tab displays the X.500 fields of the entity granting the certificate. The values apply only to available status.

Deleting an Identity Certificate

To remove an identity certificate configuration, select it, and then click **Delete**.

Note After you delete a certificate configuration, it cannot be restored. To recreate the deleted certificate, click **Add** to reenter all of the certificate configuration information.

Exporting an Identity Certificate

You can export a certificate configuration with all associated keys and certificates in PKCS12 format, which is the public key cryptography standard, and can be base64 encoded or in hexadecimal format. A complete 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 certificates 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 certificate configurations. In this case, an administrator can export a certificate configuration and then import it across the group of security appliances.

To export an identity certificate, perform the following steps:

Step 1 Click **Export** to display the Export Certificate dialog box.

Step 2 Enter the name of the PKCS12 format file to use in exporting the certificate configuration. Alternatively, click **Browse** to display the Export ID Certificate File dialog box to find the file to which you want to export the certificate configuration.

Step 3 Choose the certificate format by clicking the **PKCS12 Format** radio button or the **PEM Format** radio button.

Step 4 Enter the passphrase used to encrypt the PKCS12 file for export.

Step 5 Confirm the encryption passphrase.

Step 6 Click **Export Certificate** to export the certificate configuration.

An information dialog box appears, informing you that the certificate configuration file has been successfully exported to the location that you specified.

Generating a Certificate Signing Request

Note Entrust supports a key modulus size of 1024 *only*. Consult Entrust if you are using any other value.

To generate a certificate signing request to send to Entrust, perform the following steps:

Step 1 Click **Enroll ASA SSL VPN with Entrust** to display the Generate Certificate Signing Request dialog box.

Step 2 In the Key Pair area, perform the following steps:

a. Choose one of the configured key pairs from the drop-down list.

b. Click **Show** to display the Key Details dialog box, which provides information about the selected key pair, including date and time generated, usage (general or special purpose), modulus size, and key data.

c. Click **OK** when you are done to close Key Details dialog box.

d. Click **New** to display the Add Key Pair dialog box. To continue, go to Step 8 of the "Adding or Importing an Identity Certificate" section on page 34-9. When you generate the key pair, you can send it to the security appliance or save it to a file.

Step 3 In the Certificate Subject DN area, enter the following information:

a. The FQDN or IP address of the security appliance.

b. The name of the company.

c. The two-letter country code.

Step 4 In the Optional Parameters area, perform the following steps:

a. Click **Select** to display the Additional DN Attributes dialog box.

b. Choose the attribute to add from the drop-down list, and then enter a value.

c. Click **Add** to add each attribute to the attribute table.

d. Click **Delete** to remove an attribute from the attribute table.

e. Click **OK** when you are done to close the Additional DN Attributes dialog box.

The added attributes appear in the Additional DN Attributes field.

Step 5 Enter additional fully qualified domain name information if the CA requires it.

Step 6 Click **Generate Request** to generate the certificate signing request, which you can then send to Entrust, or save to a file and send later.

The Enroll with Entrust dialog box appears, with the CSR displayed.

Step 7 To complete the enrollment process, click the **request a certificate from Entrust** link by copying and pasting the CSR provided and submitting it through the Entrust web form, provided at http://www.entrust.net/cisco/. Alternatively, to enroll at a later time, save the generated CSR to a file, then click the **enroll with Entrust** link on the Identity Certificates pane to complete the enrollment process.

Step 8 Entrust issues a certificate after verifying the authenticity of your request. which may take several days. You then need to install the certificate by selecting the pending request in the Identity Certificate pane and clicking **Install**. Click **Close** to close the Enroll with Entrust dialog box.

Installing Identity Certificates

The Install button on the Identity Certificates pane is dimmed unless an enrollment is pending. Whenever the security appliance receives a CSR, the Identity Certificates pane displays the pending ID certificate. When you select the pending Identity Certificate, the Install button activates.

When you transmit the pending request to a CA, the CA enrolls it and returns a certificate to the security appliance. After you have received the certificate, click **Install** and highlight the appropriate identity certificate to complete the operation.

To installing a pending identity certificate, perform the following steps:

Step 1 In the Identity Certificates pane, click **Add** to display the Add Identity Certificate dialog box.

Step 2 In the Add Identity Certificate dialog box, click the **Add a new identity certificate** radio button.

Step 3 (Optional) Change the key pair or create a new key pair. A key pair is required.

Step 4 Enter the Certificate Subject DN information, and then click **Select** to display the Certificate Subject DN dialog box.

Step 5 Specify all of the subject DN attributes required by the CA involved, and then click **OK** to close the Certificate Subject DN dialog box.

Step 6 In the Add Identity Certificate dialog box, click **Advanced** to display the Advanced Options dialog box.

Step 7 To continue, see Steps 17 through 23 of the "Configuring Identity Certificates Authentication" section on page 34-8.

Step 8 In the Add Identity Certificate dialog box, click **Add Certificate**.

The Identity Certificate Request dialog box appears.

Step 9 Enter the CSR file name of type, text, such as c:\verisign-csr.txt, and then click **OK**.

Step 10 Send the CSR text file to the CA. Alternatively, you can paste the text file into the CSR enrollment page on the CA website.

Step 11 When the CA returns the Identity Certificate to you, go to the Identity Certificates pane, select the pending certificate entry, and click **Install**.

The Install Identity Certificate dialog box appears.

Step 12 Choose one of the following options by clicking the applicable radio button:

- **Install from a file**.

  Alternatively, click **Browse** to search for the file.

- **Paste the certificate data in base-64 format**.

  Paste the copied certificate data into the area provided.

Step 13 Click **Install Certificate**.

Step 14 Click **Apply** to save the newly installed certificate with the security appliance configuration.

What to Do Next

See the "Configuring Code Signer Certificates" section on page 34-14.

Configuring Code Signer Certificates

Code signing appends a digital signature to the actual executable code. This digital signature provides enough information to authenticate the signer, and ensure that the code has not been modified after being signed.

Code signer certificates are special certificates whose associated private keys are used to create digital signatures. The certificates used to sign code are obtained from a CA, in which the signed code reveals the certificate origin. You can import code signer certificates on the Code Signer pane, or choose **Configuration > Remote Access VPN > Clientless SSL VPN Access > Advanced > Java Code Signe**r.

In the Code Signer pane, you can perform the following tasks:

- Display details of a code signer certificate.
- Delete an existing code signer certificate.
- Import an existing code signer certificate.
- Export an existing code signer certificate.
- Enroll for a code signer certificate with Entrust.

This section includes the following topics:

- Showing Code Signer Certificate Details, page 34-14
- Deleting a Code Signer Certificate, page 34-15
- Importing a Code Signer Certificate, page 34-15
- Exporting a Code Signer Certificate, page 34-15

Showing Code Signer Certificate Details

To show detailed information about the selected identity certificate, click **Show Details** to display the Certificate Details dialog box, which includes the following three *display-only* tabs:

- The General tab displays the values for type, serial number, status, usage, public key type, CRL distribution point, the times within which the certificate is valid, and associated trustpoints. The values apply to both available and pending status.
- The Issued to tab displays the X.500 fields of the subject DN or certificate owner and their values. The values apply only to available status.
- The Issued by tab displays the X.500 fields of the entity granting the certificate. The values apply only to available status.

Deleting a Code Signer Certificate

To remove a code signer certificate configuration, select it, and then click **Delete**.

Note After you delete a certificate configuration, it cannot be restored. To recreate the deleted certificate, click **Import** to reenter all of the certificate configuration information.

Importing a Code Signer Certificate

To import a code signer certificate, perform the following steps:

Step 1 In the Code Signer pane, click **Import** to display the Import Certificate dialog box.

Step 2 Enter the passphrase used to decrypt the PKCS12-format file.

Step 3 Enter the name of the file to import, or click **Browse** to display the Import ID Certificate File dialog box and search for the file.

Step 4 Select the file to import and click **Import ID Certificate File**.

The selected certificate file appears in the Import Certificate dialog box.

Step 5 Click **Import Certificate**.

The imported certificate appears in the Code Signer pane.

Step 6 Click **Apply** to save the newly imported code signer certificate configuration.

Exporting a Code Signer Certificate

To export a code signer certificate, perform the following steps:

Step 1 In the Code Signer pane, click **Export** to display the Export Certificate dialog box.

Step 2 Enter the name of the PKCS12 format file to use in exporting the certificate configuration.

Step 3 In the Certificate Format area, to use the public key cryptography standard, which can be base64 encoded or in hexadecimal format, click the **PKCS12 format** radio button. Otherwise, click the **PEM format** radio button.

Step 4 Click **Browse** to display the **Export ID Certificate File** dialog box to find the file to which you want to export the certificate configuration.

Step 5 Select the file and click **Export ID Certificate File**.

The selected certificate file appears in the Export Certificate dialog box.

Step 6 Enter the passphrase used to decrypt the PKCS12 format file for export.

Step 7 Confirm the decryption passphrase.

Step 8 Click **Export Certificate** to export the certificate configuration.

What to Do Next

See the "Authenticating Using the Local CA" section on page 34-16.

Authenticating Using the Local CA

The local CA provides a secure, configurable in-house authority that resides on the security appliance for certificate authentication to use with browser-based and client-based SSL VPN connections.

Users enroll by logging in to a specified website. The local CA integrates basic certificate authority operations on the security appliance, deploys certificates, and provides secure revocation checking of issued certificates.

The local CA lets you perform the following tasks:

- Configure the local CA server.
- Revoke and unrevoke local CA certificates.
- Update CRLs.
- Add, edit, and delete local CA users.

This section includes the following topics:

- Configuring the Local CA Server, page 34-16
- Deleting the Local CA Server, page 34-19

Configuring the Local CA Server

To configure a local CA server on the security appliance, perform the following steps:

Step 1 In the CA Server pane, to activate the local CA server, click the **Enable** radio button. The default is disabled. After you enable the local CA server, the security appliance generates the local CA server certificate, key pair, and necessary database files, and then archives the local CA server certificate and key pair in a PKCS12 file.

Note Be sure to review all optional settings carefully before you enable the configured local CA. After you enable it, the certificate issuer name and key size server values cannot be changed.

The self-signed certificate key usage extension enables key encryption, key signature, CRL signature, and certificate signature.

Step 2 When you enable the local CA for the first time, you must provide an alphanumeric Enable passphrase, which must have a minimum of seven, alphanumeric characters. The passphrase protects the local CA certificate and the local CA certificate key pair archived in storage, and secures the local CA server from unauthorized or accidental shutdown. The passphrase is required to unlock the PKCS12 archive if the local CA certificate or key pair is lost and must be restored.

Note The Enable passphrase is required to enable the local CA server. Be sure to keep a record of the Enable passphrase in a safe location.

Step 3 Click **Apply** to save the local CA certificate and key pair, so the configuration is not lost if you reboot the security appliance.

Step 4 To change or reconfigure the local CA after the local CA has been configured for the first time, you must shut down the local CA server on the security appliance by clicking the **Disable** radio button. In this state, the configuration and all associated files remain in storage and enrollment is disabled.

After the configured local CA has been enabled, the following two settings are *display-only*:

- The Issuer Name field, which lists the issuer's subject name and domain name, and is formed using the username and the subject-name-default DN setting as cn=FQDN. The local CA server is the entity that grants the certificate. The default certificate name is provided in the format, cn=hostname.domainname.
- The CA Server Key Size setting, which is used for the server certificate generated for the local CA server. Key sizes can be 512, 768, 1024, or 2048 bits per key. The default is 1024 bits per key.

Step 5 From the drop-down list, choose the client key size of the key pair to be generated for each user certificate issued by the local CA server. Key sizes can be 512, 768, 1024, or 2048 bits per key. The default is 1024 bits per key.

Step 6 Enter the CA certificate lifetime value, which specifies the number of days that the CA server certificate is valid. The default is 3650 days (10 years).

The local CA server automatically generates a replacement CA certificate 30 days before expiration, which enables the replacement certificate to be exported and imported onto any other devices for local CA certificate validation of user certificates that have been issued by the local CA after they have expired.

To notify users of the upcoming expiration, the following syslog message appears in the ASDM Syslog Messages pane:

```
%ASA-1-717049: Local CA Server certificate is due to expire in days days and a replacement
certificate is available for export.
```

Note When notified of this automatic rollover, the administrator must take action to make sure that the new local CA certificate is imported to all necessary devices before it expires.

Step 7 Enter the client certificate lifetime value, which specifies the number of days that a user certificate issued by the CA server is valid. The default is 365 days (one year).

In the SMTP Server & Email Settings area, you set up e-mail access for the local CA server by specifying the following settings:

a. Enter the SMTP mail server name or IP address. Alternatively, click the ellipses (...) to display the Browse Server Name/IP Address dialog box, where you can choose the server name or IP address. Click **OK** when you are done to close the Browse Server Name/IP Address dialog box.

b. Enter the from address, from which to send e-mail messages to local CA users, in adminname@host.com format. Automatic e-mail messages carry one-time passwords to newly enrolled users and issue e-mail messages when certificates need to be renewed or updated.

c. Enter the subject, which specifies the subject line in all messages that are sent to users by the local CA server. If you do not specify a subject, the default is "Certificate Enrollment Invitation."

Step 8 To configure additional options, click the **More Options** drop-down arrow.

Step 9 Enter the CRL distribution point, which is the CRL location on the security appliance. The default location is http://hostname.domain/+CSCOCA+/asa_ca.crl.

Step 10 To make the CRL available for HTTP download on a given interface and port, choose a publish-CRL interface from the drop-down list. Then enter the port number, which can be any port number from 1-65535. The default port number is TCP port 80.

Note You cannot rename the CRL; it always has the name, LOCAL-CA-SERVER.crl.

For example, enter the URL, http://10.10.10.100/user8/my_crl_file. In this case, only the interface with the specified IP address works and when the request comes in, the security appliance matches the path, /user8/my_crl_file to the configured URL. When the path matches, the security appliance returns the stored CRL file.

Step 11 Enter the CRL lifetime in hours that the CRL is valid. The default for the CA certificate is six hours.

The local CA updates and reissues the CRL each time that a user certificate is revoked or unrevoked, but if no revocation changes occur, the CRL is reissued once every CRL lifetime. You can force an immediate CRL update and regeneration by clicking **Request CRL** in the CA Certificates pane.

Step 12 Enter the database storage location to specify a storage area for the local CA configuration and data files. The security appliance accesses and implements user information, issued certificates, and revocation lists using a local CA database. Alternatively, to specify an external file, enter the path name to the external file or click **Browse** to display the Database Storage Location dialog box.

Step 13 Choose the storage location from the list of folders that appears, and click **OK**.

Note Flash memory can store a database with 3500 users or less; a database of more than 3500 users requires external storage.

Step 14 Enter a default subject (DN string) to append to a username on issued certificates. The permitted DN attributes are provided in the following list:

- CN (Common Name)
- SN (Surname)
- O (Organization Name)
- L (Locality)
- C (Country)
- OU (Organization Unit)
- EA (E-mail Address)
- ST (State/Province)
- T (Title)

Step 15 Enter the number of hours for which an enrolled user can retrieve a PKCS12 enrollment file to enroll and retrieve a user certificate. The enrollment period is independent of the OTP expiration period. The default is 24 hours.

Note Certificate enrollment for the local CA is supported *only* for clientless SSL VPN connections. For this type of connection, communications between the client and the security appliance is through a web browser that uses standard HTML.

Step 16 Enter the length of time that a one-time password e-mailed to an enrolling user is valid. The default is 72 hours.

Step 17 Enter the number of days before expiration reminders are e-mailed to users. The default is 14 days.

Step 18 Click **Apply** to save the new or modified CA certificate configuration. Alternatively, click **Rese**t to remove any changes and return to the original settings.

Deleting the Local CA Server

To remove the local CA server from the security appliance, perform the following steps:

Step 1 In the CA Server pane, click **Delete Certificate Authority Serve**r.

The Delete Certificate Authority dialog box appears.

Step 2 To delete the CA server, click **OK**. To retain the CA server, click **Cancel**.

Note After you delete the local CA server, it cannot be restored or recovered. To recreate the deleted CA server configuration, you must reenter all of the CA server configuration information.

What to Do Next

See the "Managing the User Database" section on page 34-19.

Managing the User Database

The local CA user database includes user identification information and user status (enrolled, allowed, revoked, and so on). In the Manage User Database pane, you can perform the following tasks:

- Add a user to the local CA database.
- Change existing user identification information.
- Remove a user from the local CA database.
- Enroll a user.
- Update CRLs.
- E-mail OTPs to a user.
- View or regenerate (replace) an OTP.

This section includes the following topics:

- Adding a Local CA User, page 34-20
- Sending an Initial OTP or Replacing OTPs, page 34-20
- Editing a Local CA User, page 34-20
- Deleting a Local CA User, page 34-21
- Allowing User Enrollment, page 34-21
- Viewing or Regenerating an OTP, page 34-21

Adding a Local CA User

To add a local CA user, perform the following steps:

Step 1 To enter a new user into the local CA database, click **Add** to display the Add User dialog box.

Step 2 Enter a valid username.

Step 3 Enter an existing valid e-mail address.

Step 4 Enter the subject (DN string). Alternatively, click **Select** to display the Certificate Subject DN dialog box.

Step 5 Choose one or more DN attributes that you want to add from the drop-down list, enter a value, and then click **Add**. Available X.500 attributes for the Certificate Subject DN are the following:

- Common Name (CN)
- Department (OU)
- Company Name (O)
- Country (C)
- State/Province (ST)
- Location (L)
- E-mail Address (EA)

Step 6 Click **OK** when you are done to close the Certificate Subject DN dialog box.

Step 7 Check the **Allow enrollment** check box to enroll the user, and then click **Add User**.

The new user appears in the Manage User Database pane.

Sending an Initial OTP or Replacing OTPs

To automatically send an e-mail notice of enrollment permission with a unique OTP and the local CA enrollment URL to the newly added user, click **Email OTP**.

An Information dialog box appears indicating that the OTP was sent to the new user.

To automatically reissue a new OTP and send an e-mail notice with the new password to an existing or new user, click **Replace OTP**.

Editing a Local CA User

To modify information about an existing local CA user in the database, perform the following steps:

Step 1 Select the specific user and click **Edit** to display the Edit User dialog box.

Step 2 Enter a valid username.

Step 3 Enter an existing valid e-mail address.

Step 4 Enter the subject (DN string). Alternatively, click **Select** to display the Certificate Subject DN dialog box.

Step 5 Choose one or more DN attributes that you want to change from the drop-down list, enter a value, and then click **Add** or **Delete**. Available X.500 attributes for the Certificate Subject DN are the following:

- Common Name (CN)
- Department (OU)
- Company Name (O)
- Country (C)
- State/Province (ST)
- Location (L)
- E-mail Address (EA)

Step 6 Click **OK** when you are done to close the Certificate Subject DN dialog box.

Step 7 Check the **Allow enrollment** check box to reenroll the user, and then click **Edit User**.

The updated user details appear in the Manage User Database pane.

Deleting a Local CA User

To remove the user from the database and any certificates issued to that user from the local CA database, select the user, and then click **Delete**.

Note A deleted user cannot be restored. To recreate the deleted user record, click **Add** to reenter all of the user information.

Allowing User Enrollment

To enroll the selected user, click A**llow Enrollment**.

The status of the user changes to "enrolled" in the Manage User Database pane.

Note If the user is already enrolled, an error message appears.

Viewing or Regenerating an OTP

To view or regenerate the OTP of the selected user, perform the following steps:

Step 1 Click **View/Regenerate OTP** to display the View & Regenerate OTP dialog box.

The current OTP appears.

Step 2 After you are done, click **OK** to close the View & Regenerate OTP dialog box.

Step 3 To regenerate the OTP, click **Regenerate OTP**.

The newly generated OTP appears.

Step 4 Click **OK** to close the View & Regenerate OTP dialog box.

What to Do Next

See the “Managing User Certificates” section on page 34-22.

Managing User Certificates

To change the certificate status, perform the following steps:

Step 1 In the Manage User Certificates pane, select specific certificates by username or by certificate serial number.

Step 2 Choose one of the following options:

- If a user’s certificate lifetime period runs out, to remove that user’s access, click **Revoke**. The local CA also marks the certificate as revoked in the certificate database, automatically updates the information, and reissues the CRL.
- To restore access, select a user’s revoked certificate and click **Unrevoke**. The local CA also marks the certificate as unrevoked in the certificate database, automatically updates the certificate information, and reissues an updated CRL.

Step 3 Click **Apply** when you are done to save your changes.

What to Do Next

See the “Monitoring CRLs” section on page 34-22.

Monitoring CRLs

To monitor CRLs, perform the following steps:

Step 1 In the ASDM main application window, choose **Monitoring** > **Properties** > **CRL**.

Step 2 In the CRL area, choose the CA certificate name from the drop-down list.

Step 3 To display CRL details, click **View CRL**. For example:

```
CRL Issuer Name:
    cn=asa4.cisco.com
    LastUpdate: 09:58:34 UTC Nov 11 2009
    NextUpdate: 15:58:34 UTC Nov 11 2009
    Cached Until: 15:58:34 UTC Nov 11 2009
    Retrieved from CRL Distribution Point:
      ** CDP Not Published - Retrieved via SCEP
    Size (bytes): 224
    Associated Trustpoints: LOCAL-CA-SERVER
```

Step 4 When you are done, click **Clear CRL** to remove the CRL details and choose another CA certificate to view.

Feature History for Certificate Management

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

Table 1 Feature History for Certificate Management

Feature Name	Platform Releases	Feature Information
Certificate Management	7.0(1)	Digital certificates (including CA certificates, identity certificates, and code signer certificates) provide digital identification for authentication. A digital certificate includes information that identifies a device or user, such as the name, serial number, company, department, or IP address. CAs are trusted authorities that “sign” certificates to verify their authenticity, thereby guaranteeing the identity of the device or user. CAs issue digital certificates in the context of a PKI, which uses public-key or private-key encryption to ensure security. The following paths were introduced, based on the type of VPN connection being used: • **Configuration > Remote Access VPN > Certificate Management** • **Configuration > Site-to-Site VPN > Certificate Management**.

CHAPTER 35

Configuring IKE, Load Balancing, and NAC

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 connection.

Load balancing distributes VPN traffic among two or more security appliances in a VPN cluster.

Network Access Control (NAC) protects the enterprise network from intrusion and infection from worms, viruses, and rogue applications by performing endpoint compliance and vulnerability checks as a condition for production access to the network. We refer to these checks as *posture validation.*

This chapter describes how to configure IKE, load balancing, and NAC. It includes the following sections:

Setting IKE Parameters

This pane lets you set system wide values for VPN connections. 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 connections.

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 pane.
- Choose the second or third option for the Fragmentation Policy parameter in the Configuration > VPN > IPsec > Pre-Fragmentation pane. 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 pane. This feature is disabled by default.

This pane 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 appliances 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 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/Disables—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—Choose 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—Choose 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—Choose 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—Choose to disable aggressive mode connections.
- Alert peers before disconnecting—Choose 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—Choose 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
•	—	•	—	—

Creating 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

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—Choose 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—Choose 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—Choose the authentication method the security appliance uses to establish the identity of each IPsec peer. Preshared keys do not scale well with a growing network but are easier to set up in a small network. The choices follow:

pre-share	Preshared 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—Choose 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)	

Lifetime (secs)—Either choose 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—Choose 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
•	—	•	—	—

Assignment Policy

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.

The Assignment Policy pane lets you choose a way to assign IP addresses to remote access clients.

Fields

- Use authentication server—Choose 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 in the Configuration > AAA Setup pane.
- Use DHCP— Choose to obtain IP addresses from a DHCP server. If you use DHCP, configure the server in the Configuration > DHCP Server pane.
- Use internal address pools—Choose 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 in Configuration > Remote Access VPN > Network (Client) Access > Address Assignment > Address Pools pane.
  - Allow the reuse of an IP address __ minutes after it is released—Delays the reuse of an IP address after its return to the address pool. Adding a delay helps to prevent problems firewalls can experience when an IP address is reassigned quickly. By default, this is unchecked, meaning the security appliance does not impose a delay. If you want one, check the box and enter the number of minutes in the range 1 - 480 to delay IP address reassignment.

Modes

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

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

Address Pools

The IP Pool area 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 area 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

These panes 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 Mask—Choose 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
•	—	•	—	—

Configuring IPsec

The security appliance uses IPsec for LAN-to-LAN VPN connections, and provides the option of using IPsec for client-to-LAN VPN connections. In IPsec terminology, a “peer” is a remote-access client or another secure gateway.

Note The ASA supports LAN-to-LAN IPsec connections with Cisco peers, and with third-party peers that comply with all relevant standards.

During tunnel establishment, 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).

A LAN-to-LAN VPN connects networks in different geographic locations. 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 security 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, and 5.
- 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

Adding Crypto Maps

This pane shows the currently configured crypto maps, including the 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.

- Add—Click to launch the Create IPsec Rule dialog box, where you can configure basic, advanced, and traffic selection parameters for a rule.
- Edit—Click to edit an existing rule.
- Delete—Click to delete a rule highlighted in the table.
- Cut—Deletes a highlighted rule in the table and keeps it in the clipboard for copying.
- Copy—Copies a highlighted rule in the table.
- Find—Click to enable the Find toolbar where you can specify the parameters of existing rules that you want to find:
  - Filter—Filter the find results by selecting Interface, Source, Destination, Destination Service, or Rule Query, selecting is or contains, and entering the filter parameter. Click ... to launch a browse dialog box that displays all existing entries that you can choose.
- Diagram—Displays a diagram that illustrates the highlighted IPsec rule.
- 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.
- CA Certificate—Displays the CA certificate for the policy. This applies to static connections only.
- 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.
- Enable Anti-replay window size—Sets the anti-replay window size, between 64 and 1028 in multiples of 64. One side-effect of priority queueing in a hierarchical QoS policy with traffic shaping (see the "Rule Actions > QoS Tab") is packet re-ordering. For IPsec packets, out-of-order packets that are not within the anti-replay window generate warning syslog messages. These warnings becomes false alarms in the case of priority queueing. Configuring the anti-replay pane size helps you avoid possible false alarms.

Modes

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

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

Creating an IPsec 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 pane 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—Choose the interface name to which this policy applies.
- Policy Type—Choose the type, static or dynamic, of this tunnel policy.
- Priority—Enter the priority of the policy.
- Transform Set to Be Added—Choose 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.) Choose bidirectional, originate-only, or answer-only to specify the connection type of this policy. For LAN-to-LAN connections, choose bidirectional or answer-only (not originate-only). Choose 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 choose 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) = Use perfect forward secrecy, and use Diffie-Hellman Group 5 to generate IPsec session keys, where the prime and generator numbers are 1536 bits. This option is more secure than Group 2 but requires more processing overhead.

Modes

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

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

Creating IPsec 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.
  Reverse Route Injection (RRI) is used to populate the routing table of an internal router that runs dynanmic routing protocols such as Open Shortest Path First (OSPF), or Enhanced Interior Gateway Routing Protocol (EIGRP) , if you run ASA 8.0, or Routing Information Protocol (RIP) for remote VPN Clients or LAN²LAN sessions.
- Static Type Only Settings—Specifies parameters for static tunnel policies.
  - CA Certificate—Choose the certificate to use. If you choose 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—Chooses 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 choose Aggressive, the Diffie-Hellman Group list becomes active.
  - Diffie-Hellman Group—Choose the Diffie-Hellman group to apply. The choices are as follows: Group 1 (768-bits), Group 2 (1024-bits), or Group 5 (1536-bits).

Modes

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

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

Creating IPsec Rule/Traffic Selection Tab

This pane lets you define what traffic to protect (permit) or not protect (deny).

Fields

- Action—Specify the action for this rule to take. The selections are protect and do not protect.
- Source—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. Click ... to launch the Browse Source dialog box that contains the following fields:
  - Add/Edit—Choose IP Address or Network Object Group to add more source addresses or groups.
  - Delete—Click to delete an entry.
  - Filter—Enter an IP Address to filter the results displayed.
  - Name—Indicates that the parameters that follow specify the name of the source host or network.
  - IP Address—Indicates that the parameters that follow specify the interface, IP address, and subnet mask of the source host or network.
  - Netmask—Chooses a standard subnet mask to apply to the IP address. This parameter appears when you choose the IP Address option button.
  - Description—Enter a description.
  - Selected Source—Click **Source** to include the selected entry as a source.
- Destination—Specify the IP address, network object group or interface IP address for the destination host or network. A rule cannot use the same address as both the source and destination. Click ... to launch the Browse Destination dialog box that contains the following fields:
  - Add/Edit—Choose IP Address or Network Object Group to add more destination addresses or groups.
  - Delete—Click to delete an entry.
  - Filter—Enter an IP Address to filter the results displayed.
  - Name—Indicates that the parameters that follow specify the name of the destination host or network.
  - IP Address—Indicates that the parameters that follow specify the interface, IP address, and subnet mask of the destination host or network.
  - Netmask—Chooses a standard subnet mask to apply to the IP address. This parameter appears when you choose the IP Address option button.
  - Description—Enter a description.
  - Selected Destination—Click **Destination** to include the selected entry as a destination.
- Service—Enter a service or click ... to launch the browse service dialog box where you can choose from a list of services.
- Description—Enter a description for the Traffic Selection entry.
- More Options
  - Enable Rule—Click to enable this rule.
  - Source Service—Enter a service or click **...** to launch the browse service dialog box where you can choose from a list of services.
  - Time Range—Define a time range for which this rule applies.

  - Group—Indicates that the parameters that follow specify the interface and group name of the source host or network.
  - Interface—Choose the interface name for the IP address. This parameter appears when you choose the IP Address option button.
  - IP address—Specifies the IP address of the interface to which this policy applies. This parameter appears when you choose the IP Address option button.
  - Destination—Specify the IP address, network object group or interface IP address for the source or destination host or network. A rule cannot use the same address as both the source and destination. Click ... for either of these fields to launch the Browse dialog box that contain the following fields:
  - Name—Choose the interface name to use as the source or destination host or network. This parameter appears when you choose the Name option button. This is the only parameter associated with this option.
  - Interface—Choose the interface name for the IP address. This parameter appears when you choose the Group option button.
  - Group—Choose 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 choose the Group option button.
- **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 pane, 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 chose 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 choose 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)—Choose 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

Use this pane 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

Use this pane to modify an existing IPsec pre-fragmentation policy and do-not-fragment (DF) bit policy for an interface selected on the parent pane, **Configuration > VPN > IPsec > Pre-Fragmentation**

Fields

- **Interface**—Identifies the chosen 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**—Choose 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
•	—	•	—	—

IPsec Transform Sets

Use this pane 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

Use this pane 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**—Choose 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**—Choose 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
•	—	•	—	—

Configuring Load Balancing

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 To use VPN load balancing, you must have an ASA Model 5510 with a Plus license or an ASA Model 5520 or higher. VPN load balancing also 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.

The following sections explain load balancing:

- Eligible Clients
- Enabling Load Balancing

- Creating Virtual Clusters
- Mixed Cluster Scenarios
- Comparing Load Balancing to Failover
- Load Balancing Prerequisites

Eligible Clients

Load balancing is effective only on remote sessions initiated with the following clients:

- Cisco AnyConnect VPN Client (Release 2.0 and later)
- Cisco VPN Client (Release 3.0 and later)
- Cisco ASA 5505 Security Appliance (when acting as an Easy VPN client)
- Cisco VPN 3002 Hardware Client (Release 3.5 or later)
- Cisco PIX 501/506E when acting as an Easy VPN client
- IOS EZVPN Client devices supporting IKE-redirect (IOS 831/871)
- Clientless SSL VPN (not a client)

Load balancing works with IPsec clients and SSL VPN client and clientless sessions. All other VPN connection types (L2TP, PPTP, L2TP/IPsec), including LAN-to-LAN, can connect to an security appliance on which load balancing is enabled, but they cannot participate in load balancing.

Enabling Load Balancing

This pane 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.

Creating Virtual Clusters

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 *backup 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 backup 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 backup 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

A load-balancing cluster can consist of security appliances of the same release, of mixed releases, as well as VPN 3000 concentrators, or a mixture of these, subject to the following restrictions:

- Load-balancing clusters that consist of both same release security appliances and VPN 3000 concentrators can run load balancing for a mixture of IPsec, AnyConnect, and clientless SSL VPN client and clientless sessions.
- Load-balancing clusters that include mixed release security appliances or same release security appliances and VPN 3000 concentrators or both can support only IPsec sessions. In such a configuration, however, the security appliances might not reach their full IPsec capacity. "Scenario 1: Mixed Cluster with No SSL VPN Connections" on page 22, illustrates this situation.

Since Release 7.1(1), IPsec and SSL VPN sessions count or weigh equally in determining the load that each device in the cluster carries. This represents a departure from the load balancing calculation for the ASA Release 7.0(x) software and the VPN 3000 concentrator, in that these platforms both use a weighting algorithm that, on some hardware platforms, calculates SSL VPN session load differently from IPsec session load.

The virtual master of the cluster assigns session requests to the members of the cluster. The security appliance regards all sessions, SSL VPN or IPsec, as equal and assigns them accordingly. You can configure the number of IPsec and SSL VPN sessions to allow, up to the maximum allowed by your configuration and license.

We have tested up to ten nodes in a load-balancing cluster. Larger clusters might work, but we do not officially support such topologies.

Mixed Cluster Scenarios

If you have a mixed configuration—that is, if your load-balancing cluster includes devices running a mixture of ASA software releases or at least one security appliance running ASA Release 7.1(1) or later and a VPN 3000 concentrator—the difference in weighting algorithms becomes an issue if the initial cluster master fails and another device takes over as master.

The following scenarios illustrate the use of VPN load balancing in clusters consisting of a mixture of security appliances running ASA Release 7.1(1) and ASA Release 7.0(x) software, as well as VPN 3000 Series Concentrators.

Scenario 1: Mixed Cluster with No SSL VPN Connections

In this scenario, the cluster consists of a mixture of security appliances and VPN 3000 Concentrators. Some of the security appliance cluster peers are running ASA Release 7.0(x), and some are running Release 7.1(1). The pre-7.1(1) and VPN 3000 peers do not have any SSL VPN connections, and the 7.1(1) cluster peers have only the base SSL VPN license, which allows two SSL VPN sessions, but there are no SSL VPN connections. In this case, all the connections are IPsec, and load balancing works fine.

The two SSL VPN licenses have a very small effect on the user's taking advantage of the maximum IPsec session limit, and then only when a VPN 3000 Concentrator is the cluster master. In general, the smaller the number of SSL VPN licenses is on a security appliance in a mixed cluster, the smaller the effect on the ASA 7.1(1) device being able to reach its IPsec session limit in a scenario where there are only IPsec sessions.

Scenario 2: Mixed Cluster Handling SSL VPN Connections

Suppose, for example, a security appliance running ASA Release 7.1(1) software is the initial cluster master; then that device fails. Another device in the cluster takes over automatically as master and applies its own load-balancing algorithm to determine processor loads within the cluster. A cluster master running ASA Release 7.1(1) software cannot weight session loads in any way other than what that software provides. Therefore, it cannot assign a combination of IPsec and SSL VPN session loads properly to ASA devices running earlier versions nor to VPN 3000 Concentrators. Conversely, a VPN 3000 Concentrator acting as the cluster master cannot assign loads properly to an ASA Release 7.1(1) security appliance. The following scenario illustrates this dilemma.

This scenario is similar to the previous one, in that the cluster consists of a mixture of security appliances and VPN 3000 Concentrators. Some of the security appliance cluster peers are running ASA Release 7.0,(x) and some are running Release 7.1(1). In this case, however, the cluster is handling SSL VPN connections as well as IPsec connections.

If a device that is running software earlier than ASA Release 7.1(1) is the cluster master, the master applies the protocol and logic in effect prior to Release 7.1(1). That is, sessions might be directed to load-balancing peers that have exceeded their session limit. In that case, the user is denied access.

If the cluster master is a device running ASA Release 7.0(x) software, the old session-weighting algorithm applies only to the pre-7.1(1) peers in the cluster. No one should be denied access in this case. Because the pre-7.1(1) peers use the session-weighting algorithm, they are more lightly loaded.

An issue arises, however, because you cannot guarantee that the 7.1(1) peer is always the cluster master. If the cluster master fails, another peer assumes the role of master. The new master might be any of the eligible peers. Because of the innately unpredictability of the results, we recommend that you avoid configuring this type of cluster.

Comparing Load Balancing to Failover

Both load balancing and failover are high-availability features, but they function differently and have different requirements. In some circumstances you can use both load balancing and failover. The following sections describe the differences between these features.

Load balancing is a mechanism for equitably distributing remote-access VPN traffic among the devices in a virtual cluster. It is based on simple distribution of traffic without taking into account throughput or other factors. A load-balancing cluster consists of two or more devices, one of which is the virtual master, and the others backup. These devices do not need to be of the exact same type, or have identical software versions or configurations. All active devices in a virtual cluster carry session loads. Load balancing directs traffic to the least loaded device in the cluster, distributing the load among all devices. It makes efficient use of system resources and provides increased performance and high availability.

A *failover* configuration requires two identical security appliances connected to each other through a dedicated failover link and, optionally, a stateful failover link. The health of the active interfaces and units is monitored to determine when specific failover conditions are met. If those conditions occur, failover occurs. Failover supports both VPN and firewall configurations.

The security appliance supports two failover configurations, Active/Active failover and Active/Standby failover. VPN connections run only in Active/Standby, single routed mode. Active/Active failover requires multi-context mode, so does not support VPN connections.

With Active/Active failover, both units can pass network traffic. This is not true load balancing, although it might appear to have the same effect. When failover occurs, the remaining active unit takes over passing the combined traffic, based on he configured parameters. Therefore, when configuring Active/Active failover, you must make sure that the combined traffic for both units is within the capacity of each unit.

With Active/Standby failover, only one unit passes traffic, while the other unit waits in a standby state and does not pass traffic. Active/Standby failover lets you use a second security applianceto take over the functions of a failed unit. When the active unit fails, it changes to the standby state, while the standby unit changes to the active state. The unit that becomes active assumes the IP addresses (or, for transparent firewall, the management IP address) and MAC addresses of the failed unit and begins passing traffic. The unit that is now in standby state takes over the standby IP addresses of the active unit. If an active unit fails, the standby takes over without any interruption to the client VPN tunnel.

Load Balancing 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. All of the outside and inside network interfaces on the load-balancing devices in a cluster must be on the same IP network.

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 check this 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, check this 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 check 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 backup 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. If NAT is not being used (or if the device is not behind a firewall using NAT), leave the field blank.
  - **Send FQDN to client**—Check this check box to cause the VPN cluster master to send a fully qualified domain name using the host and domain name of the cluster device instead of the outside IP address when redirecting VPN client connections to that cluster device.

    By default, the ASA sends only IP addresses in load-balancing redirection to a client. If certificates are in use that are based on DNS names, the certificates will be invalid when redirected to a backup device.

    As a VPN cluster master, this security appliance can send a fully qualified domain name (FQDN), using reverse DNS lookup, of a cluster device (another security appliance in the cluster), instead of its outside IP address, when redirecting VPN client connections to that cluster device.

    All of the outside and inside network interfaces on the load-balancing devices in a cluster must be on the same IP network.

To enable Clientless SSL VPN load balancing using FQDNs rather than IP addresses, perform the following configuration steps:

Step 1 Enable the use of FQDNs for Load Balancing by checking the Send FQDN to client... checkbox.

Step 2 Add an entry for each of your security appliance outside interfaces into your DNS server, if such entries are not already present. Each security appliance outside IP address should have a DNS entry associated with it for lookups. These DNS entries must also be enabled for Reverse Lookup.

Step 3 Enable DNS lookups on your security appliance on the dialog box Configuration > Device Management > DNS > DNS Client for whichever interface has a route to your DNS server.

Step 4 Define your DNS server IP address on the security appliance. To do this, click Add on this dialog box. This opens the Add DNS Server Group dialog box. Enter the IP address of the DNS server you want to add; for example, `192.168.1.1` (IP address of your DNS server).

Step 5 Click OK and Apply.

Modes

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

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

Setting Global NAC Parameters

The security appliance uses Extensible Authentication Protocol (EAP) over UDP (EAPoUDP) messaging to validate the posture of remote hosts. Posture validation involves checking 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.

Fields

The NAC pane lets you set attributes that apply to all NAC communications. The following global attributes at the top of the pane apply to EAPoUDP messaging between the security appliance and remote hosts:

- 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.
- Retry if no response—Number of times the security appliance resends an EAP over UDP message. This attribute limits the number of consecutive retries sent in response to Rechallenge Interval expirations. The setting is in seconds. Enter a value in the range 1 to 3. The default setting is 3.
- Rechallenge Interval—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.

- Wait before new PV Session—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 value of the "Retry if no response" setting. 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.

The Clientless Authentication area of the NAC pane 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—Click 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 leave this attribute blank, the security appliance applies the default ACL for clientless hosts.
- Clientless 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.
- Enable Audit—Click to pass the IP address of the client to an optional audit server if the client does not respond to a posture validation request. The audit server, such as a Trend server, uses the host IP address to challenge the host directly to assess its health. For example, it may challenge the host to determine whether its virus checking software is active and up-to-date. After the audit server completes its interaction with the remote host, it passes a token to the posture validation server, indicating the health of the remote host.
- None—Click to disable clientless authentication and audit services.

Modes

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

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

Configuring Network Admission Control Policies

The NAC Policies table displays the Network Admission Control (NAC) policies configured on the security appliance.

To add, change, or remove a NAC policy, do one of the following:

- To add a NAC policy, choose **Add**. The Add NAC Framework Policy dialog box opens.
- To change a NAC policy, double-click it, or select it and click **Edit**. The Edit NAC Framework Policy dialog box opens.

- To remove a NAC policy, select it and click **Delete**.

The following sections describe NAC, its requirements, and how to assign values to the policy attributes:

- About NAC
- Uses, Requirements, and Limitations
- Fields
- What to Do Next

About NAC

NAC protects the enterprise network from intrusion and infection from worms, viruses, and rogue applications by performing endpoint compliance and vulnerability checks as a condition for production access to the network. We refer to these checks as *posture validation.* You can configure posture validation to ensure that the anti-virus files, personal firewall rules, or intrusion protection software on a host with an AnyConnect or Clientless SSL VPN session are up-to-date before providing access to vulnerable hosts on the intranet. Posture validation can include the verification that the applications running on the remote hosts are updated with the latest patches. NAC occurs only after user authentication and the setup of the tunnel. NAC is especially useful for protecting the enterprise network from hosts that are not subject to automatic network policy enforcement, such as home PCs.

The establishment of a tunnel between the endpoint and the security appliance triggers posture validation.

You can configure the security appliance to pass the IP address of the client to an optional audit server if the client does not respond to a posture validation request. The audit server, such as a Trend server, uses the host IP address to challenge the host directly to assess its health. For example, it may challenge the host to determine whether its virus checking software is active and up-to-date. After the audit server completes its interaction with the remote host, it passes a token to the posture validation server, indicating the health of the remote host.

Following successful posture validation or the reception of a token indicating the remote host is healthy, the posture validation server sends a network access policy to the security appliance for application to the traffic on the tunnel.

In a *NAC Framework* configuration involving the security appliance, only a Cisco Trust Agent running on the client can fulfill the role of posture agent, and only a Cisco Access Control Server (ACS) can fulfill the role of posture validation server. The ACS uses dynamic ACLs to determine the access policy for each client.

As a RADIUS server, the ACS can authenticate the login credentials required to establish a tunnel, in addition to fulfilling its role as posture validation server.

Note Only a NAC Framework policy configured on the security appliance supports the use of an audit server.

In its role as posture validation server, the ACS uses access control lists. If posture validation succeeds and the ACS specifies a redirect URL as part of the access policy it sends to the security appliance, the security appliance redirects all HTTP and HTTPS requests from the remote host to the redirect URL. Once the posture validation server uploads an access policy to the security appliance, all of the associated traffic must pass both the Security Appliance and the ACS (or vice versa) to reach its destination.

The establishment of a tunnel between a remote host and the security appliance triggers posture validation if a NAC Framework policy is assigned to the group policy. The NAC Framework policy can, however, identify operating systems that are exempt from posture validation and specify an optional ACL to filter such traffic.

Uses, Requirements, and Limitations

When configured to support NAC, the security appliance functions as a client of a Cisco Secure Access Control Server, requiring that you install a minimum of one Access Control Server on the network to provide NAC authentication services.

Following the configuration of one or more Access Control Servers on the network, you must register the Access Control Server group, using the **Configuration > Remote Access VPN > Clientless SSL VPN Access > Group Policies > Add or Edit External** menu option. Then add the NAC policy.

ASA support for NAC Framework is limited to remote access IPsec and Clientless SSL VPN sessions. The NAC Framework configuration supports only single mode.

NAC on the ASA does not support Layer 3 (non-VPN) and IPv6 traffic.

Fields

- Policy Name—Enter a string of up to 64 characters to name the new NAC policy.

  Following the configuration of the NAC policy, the policy name appears next to the NAC Policy attribute in the Network (Client) Access group policies. Assign a name that will help you to distinguish its attributes or purpose from others that you may configure.
- Status Query Period—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 Period—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.
- Authentication Server Group—Specifies the authentication server group to use for posture validation. The drop-down list next to this attribute displays the names of all server groups of type RADIUS configured on this security appliance that are available for remote access tunnels. Select an ACS group consisting of at least one server configured to support NAC.
- 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.

What to Do Next

Following the configuration of the NAC policy, you must assign it to a group policy for it to become active. To do so, choose **Configuration > Remote Access VPN> Network (Client) Access > Group Policies > Add** or **Edit > General** > **More Options** and the NAC policy name from the drop-down list next to the NAC Policy attribute.

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
•	—	•	—	—

Add/Edit Posture Validation Exception

The Add/Edit Posture Validation Exception dialog pane 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 pane 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
•	—	•	—	—

CHAPTER 36

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.

Client Software

The Client Software pane lets administrators at a central location do the following actions:

- Enable client 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 in Configuration > Remote Access VPN > Network (Client) Access > Advanced > IPsec > Upload Software > Client Software applies only to the IPsec VPN client, (For Windows, MAC OS X, and Linux), and the VPN 3002 hardware client. It does not apply to the Cisco AnyConnect VPN clients, which is updated by the security appliance automatically when it connects.

For the IPsec VPN client, 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 Site-to-Site IPsec connection or a Clientless VPN IPsec connection, you do not receive an error message, but no update notification or client update goes to those types of IPsec connections.

To enable client update globally for all clients of a particular client type, use this dialog box. You can also notify all Windows, MAC OS X, and Linux clients that an upgrade is needed and initiate an upgrade on all VPN 3002 hardware clients from this dialog box. 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, choose **Configuration > Remote Access VPN > Network (Client) Access > IPsec > Add/Edit > Advanced > IPsec > Client Software Update**.

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, MAC OS X, and Linux VPN clients, or initiate an automatic update to hardware clients.
- Client Type—Lists the clients to upgrade: software or hardware, and for Windows software clients, all Windows 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 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, MAC OS X, and Linux-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, MAC OS X, and Linux-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 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** 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

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 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, MAC OS X, or Linux-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, MAC OS X, or Linux-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

To configure the default tunnel gateway, click the **Static Route** link. The Configuration > Routing > Routing > Static Route dialog box 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 Policies

The Group Policies pane 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" panes and dialog boxes let you configure the group parameters, including those for the default group. The default group 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.

You can configure either an internal or an external group policy. An internal group policy is stored locally, and an external group policy is stored externally on a RADIUS or LDAP server. Clicking Edit opens a similar dialog box on which you can create a new group policy or modify an existing one.

In these dialog boxes, you configure the following kinds of parameters:

- General attributes: Name, banner, address pools, protocols, filtering, and connection settings.
- Servers: DNS and WINS servers, DHCP scope, and default domain name.
- Advanced attributes: Split tunneling, IE browser proxy, SSL VPN client and AnyConnect client, and IPsec client.

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—Lists 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—Offers a drop-down menu on which you can select whether to add an internal or an external group policy. If you simply click Add, then by default, you create an internal group policy. Clicking Add opens the Add Internal Group Policy dialog box or the Add External Group Policy dialog box, which let you add a new group policy to the list. This dialog box includes three menu sections. Click each menu item to display its parameters. As you move from item to item, ASDM retains your settings. When you have finished setting parameters on all menu sections, click **Apply** or **Cancel**. Offers a drop-down menu from which you can select whether to add an internal or an external group policy. If you simply click Add, then by default, you create an internal group policy.

  - Edit—Displays the Edit Group Policy dialog box, which lets you modify an existing 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

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

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.

Adding or Editing a Remote Access Internal Group Policy, General Attributes

The Add or Edit Group Policy dialog box 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 dialog box, checking the Inherit check box lets the corresponding setting take its value from the default group policy. Inherit is the default value for all of the attributes in this dialog box.

Fields

The following attributes appear in the Add Internal Group Policy > General dialog box. They apply to SSL VPN and IPsec sessions, or clientless SSL VPN sessions. Thus, several are present for one type of session, but not the other.

- Name—Specifies the name of this group policy up to 64 characters; spaces are allowed. For the Edit function, this field is read-only.
- Banner—Specifies the banner text to present to users at login. The length can be up to 491 characters. There is no default value.
- Address Pools—(Network (Client) Access only) Specifies the name of one or more address pools to use for this group policy.
- Select—(Network (Client) Access only) Opens the Select Address Pools dialog box, which shows the pool name, starting and ending addresses, and subnet mask of address pools available for client address assignment and lets you select, add, edit, delete, and assign entries from that list.
- IPv6 Address Pools—Specifies the name of one or more IPv6 address pools to use for this group policy. The Select button following this field opens the Select Address Pools dialog box, as previously described.
- More Options—Displays additional configurable options for this group policy.
- Tunneling Protocols—Specifies the tunneling protocols that this group can use. Users can use only the selected protocols. The choices are as follows:
  - Clientless SSL VPN—Specifies the use of VPN via SSL/TLS, which uses a web browser to establish a secure remote-access tunnel to a security appliance; requires neither a software nor hardware client. Clientless SSL VPN 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.
  - SSL VPN Client—Specifies the use of the Cisco AnyConnect VPN client or the legacy SSL VPN client.

- IPsec—IP Security Protocol. Regarded as the most secure protocol, IPsec provides the most complete architecture for VPN tunnels. Both Site-to-Site (peer-to-peer) connections and client-to-LAN connections can use IPsec.
- 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 you do not select a protocol, an error message appears.

- IPv4Filter—(Network (Client) Access only) Specifies which access control list to use for an IPv4 connection, 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 ACL Manager dialog box.
- IPv6Filter—(Network (Client) Access only) Specifies which access control list to use for an IPv6 connection, 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 ACL Manager dialog box.
- Web ACL—(Clientless SSL VPN only) Choose an access control list (ACL) from the drop-down list if you want to filter traffic. Click Manage next to the list if you want to view, modify, add, or remove ACLs before making a selection.
- Manage—Displays the ACL Manager dialog box, 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 dialog box.
- NAC Policy—Selects the name of a Network Admission Control policy to apply to this group policy. You can assign an optional NAC policy to each group policy. The default value is --None--.
- Manage—Opens the Configure NAC Policy dialog box. After configuring one or more NAC policies, the NAC policy names appear as options in the drop-down list next to the NAC Policy attribute.
- Access Hours—Selects 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 checked, the default value is --Unrestricted--.
- Manage—Opens the Browse Time Range dialog box, in which you can add, edit, or delete a time range.
- Simultaneous Logins—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.

- Restrict Access to VLAN—(Optional) Also called “VLAN mapping,” this parameter specifies the egress VLAN interface for sessions to which this group policy applies. The security appliance forwards all traffic on this group to the selected VLAN. Use this attribute to assign a VLAN to the group policy to simplify access control. Assigning a value to this attribute is an alternative to using ACLs to filter traffic on a session. In addition to the default value (Unrestricted), the drop-down list shows only the VLANs that are configured on this security appliance.

Note This feature works for HTTP connections, but not for FTP and CIFS.

- Maximum Connect Time—If the Inherit check box is not checked, 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, check Unlimited (the default).
- Idle Timeout—If the Inherit check box is not checked, this parameter specifies this user's idle timeout period in minutes. If there is no communication activity on the user 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, check **Unlimited**. This value does not apply to Clientless SSL VPN users.
- On smart card removal—With the default option, Disconnect, the client tears down the connection if the smart card used for authentication is removed. Click **Keep the connection** if you do not want to require users to keep their smart cards in the computer for the duration 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
•	—	•	—	—

Configuring the Portal for a Group Policy

The Portal attributes determine what appears on the portal page for members of this group policy establishing Clientless SSL VPN connections. In this pane, you can enable Bookmark lists and URL Entry, file server access, Port Forwarding and Smart Tunnels, ActiveX Relay, and HTTP settings.

Fields

- Bookmark List—Choose a previously-configured Bookmark list or click **Manage** to create a new one. Bookmarks appear as links, from which users can navigate from the portal page.
- URL Entry—Enable to allow remote users to enter URLs directly into the portal URL field.
- File Access Control—Controls the visibility of "hidden shares" for Common Internet File System (CIFS) files. A hidden share is identified by a dollar sign ($) at the end of the share name. For example, drive C is shared as C$. With hidden shares, a shared folder is not displayed, and users are restricted from browsing or accessing these hidden resources.
  - File Server Entry—Enable to allow remote users to enter the name of a file server.
  - File Server Browsing—Enable to allow remote users to browse for available file servers.
  - Hidden Share Access—Enable to hide shared folders.

- Port Forwarding Control—Provides users access to TCP-based applications over a Clientless SSL VPN connection through a Java Applet.
  - Port Forwarding List—Choose a previously-configured list TCP applications to associate with this group policy. Click **Manage** to create a new list or to edit an existing list.
  - Auto Applet Download—Enables automatic installation and starting of the Applet the first time the user logs in.
  - Applet Name—Changes the name of the title bar that of the Applet dialog box to the name you designate. By default, the name is Application Access.
- Smart Tunnel—Connects a Winsock 2, TCP-based application installed on the end station to a server on the intranet, using a clientless (browser-based) SSL VPN session with the security appliance as the pathway, and the security appliance as a proxy server.
  - Smart Tunnel List—Choose the list name from the drop-down menu if you want to provide smart tunnel access. Assigning a smart tunnel list to a group policy or username enables smart tunnel access for all users whose sessions are associated with the group policy or username, but restricts smart tunnel access to the applications specified in the list. To view, add, modify, or delete a smart tunnel list, click **Manage**.
  - Auto Start (Smart Tunnel List)—Check the check box to start smart tunnel access automatically upon user login. Uncheck the check box to enable smart tunnel access upon user login, but require the user to start it manually, using the Application Access > Start Smart Tunnels button on the Clientless SSL VPN Portal Page.
  - Auto Sign-on Server List—Choose the list name from the drop-down menu if you want to reissue the user credentials when the user establishes a smart tunnel connection to a server. Each smart tunnel auto sign-on list entry identifies a server with which to automate the submission of user credentials. To view, add, modify, or delete a smart tunnel auto sign-on list, click **Manage**.
  - Windows Domain Name (Optional)—Specify the Windows domain to add it to the username during auto sign-on, if the universal naming convention (domain\username) is required for authentication. For example, enter CISCO to specify CISCO\qa_team when authenticating for the username qu_team. You must also check the “Use Windows domain name with user name” option when configuring associated entries in the auto sign-on server list.
- ActiveX Relay—Lets Clientless users launch Microsoft Office applications from the browser. The applications use the session to download and upload Microsoft Office documents. The ActiveX relay remains in force until the Clientless SSL VPN session closes.

More Options:

- HTTP Proxy—Enables or disables 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 automatically modifies the old browser proxy configuration 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.
- Auto Start (HTTP Proxy)—Check to enable HTTP Proxy automatically upon user login. Uncheck to enable smart tunnel access upon user login, but require the user to start it manually.
- HTTP Compression—Enables compression of HTTP data over the Clientless SSL VPN session.

Modes

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

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

Configuring Customization for a Group Policy

To configure customization for a group policy, select a preconfigured portal customization object, or accept the customization provided in the default group policy. You can also configure a URL to display

Fields

- Portal Customization—Configure a customization object for the end user portal.
  - Inherit—To inherit a portal customization from the default group policy, check **Inherit**. To specify a previously configured customization object, uncheck Inherit and choose the customization object from the drop-down list.
  - Manage—Click to import a new customization object.
- Homepage URL (optional)—To specify a homepage URL for users associated with the group policy, enter it in this field. The string must begin with either http:// or https://. To inherit a home page from the default group policy, click **Inherit**. Clientless users are immediately brought to this page after successful authentication. AnyConnect launches the default web browser to this URL upon successful establishment of the VPN connection. On Linux platforms, AnyConnect does not currently support this field and ignores it.
- Access Deny Message—To create a message to users for whom access is denied, enter it in this field. To accept the message in the default group policy, click **Inherit**.

Modes

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

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

Adding or Editing a Site-to-Site Internal Group Policy

The Add or Edit Group Policy dialog box lets you specify tunneling protocols, filters, connection settings, and servers for the group policy being added or modified. For each of the fields in this dialog box, checking the Inherit check box lets the corresponding setting take its value from the default group policy. Inherit is the default value for all of the attributes on this dialog box.

Fields

The following attributes appear in the Add Internal Group Policy > General dialog box. They apply to SSL VPN and IPsec sessions, or clientless SSL VPN sessions. Thus, several are present for one type of session, but not the other.

- Name—Specifies the name of this group policy. For the Edit function, this field is read-only.
- Tunneling Protocols—Specifies the tunneling protocols that this group can use. Users can use only the selected protocols. The choices are as follows:
  - Clientless SSL VPN—Specifies the use of VPN via SSL/TLS, which uses a web browser to establish a secure remote-access tunnel to a security appliance; requires neither a software nor hardware client. Clientless SSL VPN 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.
  - SSL VPN Client—Specifies the use of the Cisco AnyConnect VPN client or the legacy SSL VPN client.
  - IPsec—IP Security Protocol. Regarded as the most secure protocol, IPsec provides the most complete architecture for VPN tunnels. Both Site-to-Site (peer-to-peer) connections and client-to-LAN connections can use IPsec.
  - L2TP/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 you do not select a protocol, an error message appears.

- Filter—(Network (Client) Access only) Specifies which access control list 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 Group Policy dialog box.
- Manage—Displays the ACL Manager dialog box, 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 dialog box.

Browse Time Range

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, in 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, in 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

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, in which you can configure a recurring time range.
- Edit—Opens the Edit Recurring Time Range dialog box, in 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

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—Specifies 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 check 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

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 section.

Standard ACL

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, in 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, in 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

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, in 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, in 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 click Edit Log Option. The Log Options dialog box 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

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. Both IPv4 and IPv6 addresses are supported. With this selection, the IP Address, ellipsis button, and Netmask fields become available. Choose 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. Choose 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 choose 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 choose a protocol from a drop-down list, choose 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

The Browse Source or Destination Address dialog box lets you select an object to use as 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 choose 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

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, in 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

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

The Browse ICMP dialog box lets you select an ICMP group for this rule.

Fields

- Add—Opens the Add ICMP Group dialog box, in 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

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. Choosing 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

The Browse Other dialog box lets you select a protocol group for this rule.

Fields

- Add—Opens the Add Protocol Group dialog box, in 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

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 > Servers

The Add or Edit Group Policy dialog box, Servers item lets you specify DNS and WINS servers, as well as the DHCP scope and default domain.

Login Setting

In this dialog box, you can enable the security appliance to prompt remote users to download the AnyConnect client. Figure 36-1 shows the prompt displayed:

Figure 36-1 Prompt Displayed to Remote Users for SSL VPN Client Download

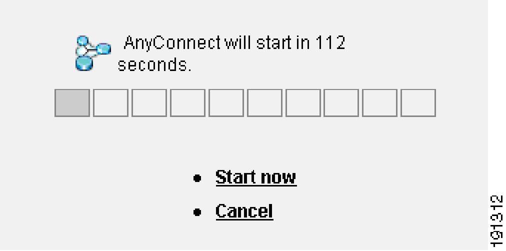

AnyConnect will start in 112 seconds.

- Start now
- Cancel

Fields

- Post Login Setting—Choose to prompt the user and set the timeout to perform the default post login selection.
- Default Post Login Selection—Choose an action to perform after login.

Modes

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

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

Key Regeneration

Rekey Negotiation occurs when the security appliance and the client perform a rekey and they renegotiate the crypto keys and initialization vectors, increasing the security of the connection.

Fields

- Renegotiation Interval—Uncheck the Unlimited check box to specify the number of minutes from the start of the session until the rekey takes place, from 1 to 10080 (1 week).
- Renegotiation Method—Check the **None** check box to disable rekey, check the **SSL** check box to specify SSL renegotiation during a rekey, or check the **New Tunnel** check box to establish a new tunnel during rekey.

Modes

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

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

Dead Peer Detection

Dead Peer Detection (DPD) ensures that the security appliance (gateway) or the client can quickly detect a condition where the peer is not responding, and the connection has failed.

Fields

- Gateway Side Detection—Uncheck the **Disable** check box to specify that DPD is performed by the security appliance (gateway). Enter the interval, from 30 to 3600 seconds, with which the security appliance performs DPD.
- Client Side Detection—Uncheck the **Disable** check box to specify that DPD is performed by the client. Enter the interval, from 30 to 3600 seconds, with which the client performs DPD.

Modes

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

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

Customization

Fields

- Portal Customization—Selects the customization to apply to the AnyConnect Client/SSL VPN portal page. The default is DfltCustomization.
- Manage—Opens the Configure GUI Customization objects dialog box, in which you can specify that you want to add, edit, delete, import, or export a customization object.
- Access Deny Message—Specifies a message to display to the end user when the connection is denied. Select Inherit to accept the message in the default group policy. The default message, if you deselect Inherit, is: “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
•	—	•	—	—

ACLs

This dialog box lets you configure ACLs for Clientless SSL VPN.

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 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
•	—	•	—	—

Group Policy Advanced SSL VPN Client

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 pane.
- Keep Installer on Client System—Enable to allow permanent client installation on the remote computer. Enabling disables the automatic uninstalling feature of the client. The client remains installed on the remote computer for subsequent connections, reducing the connection time for the remote user.
- Compression—Compression increases the communications performance between the security appliance and the client by reducing the size of the packets being transferred.

- Datagram Transport Layer Security (DTLS)—DTLS avoids latency and bandwidth problems associated with some SSL connections and improves the performance of real-time applications that are sensitive to packet delays.
- Keepalive Messages—Enter an number, from 15 to 600 seconds, in the Interval field to enable and adjust the interval of keepalive messages to ensure that an 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 interval also ensures that the client 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.
- MTU—Adjusts the MTU size for SSL connections. Enter a value in bytes, from 256 to 1410 bytes. By default, the MTU size is adjusted automatically based on the MTU of the interface that the connection uses, minus the IP/UDP/DTLS overhead.
- Client Profile to Download—a profile is a group of configuration parameters that the AnyConnect client uses to configure the connection entries that appear in the user interface, including the names and addresses of host computers.
- Optional Client Module to Download—To minimize download time, the AnyConnect client only requests downloads (from the security appliance) of modules that it needs for each feature that it supports. You must specify the names of modules that enable other features, such as *sbl* to enable the feature Start Before Logon (SBL). Separate module names in the string with a comma.

  For a list of values to enter for each client feature, see the Cisco AnyConnect VPN Client Administrator Guide.

Modes

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

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

Login Setting

In this dialog box, you can enable the security appliance to prompt remote users to download the AnyConnect client. Figure 36-1 shows the prompt displayed:

Figure 36-2 Prompt Displayed to Remote Users for SSL VPN Client Download

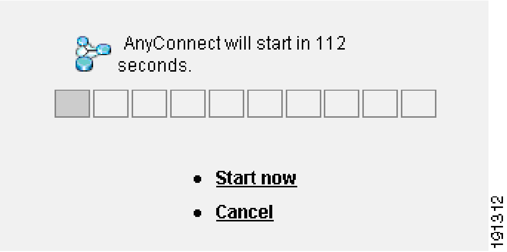

AnyConnect will start in 112 seconds.

- Start now
- Cancel

Fields

- Inherit—Check the Inherit check box to inherit the value from the default group policy.
- Post Login Setting—Click to indicate whether to prompt the user to choose the post-login action and, if you choose to prompt, set the timeout to perform the default post login selection.
- Default Post Login Selection—Click an action to perform after login: either go to the clientless SSL VPN portal or download the SSL VPN (AnyConnect) client.

Modes

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

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

Key Regeneration

Rekey Negotiation occurs when the security appliance and the client perform a rekey and they renegotiate the crypto keys and initialization vectors, increasing the security of the connection.

Fields

- Renegotiation Interval**—**Clear the **Unlimited** check box to specify the number of minutes from the start of the session until the rekey takes place, from 1 to 10080 (1 week).
- Renegotiation Method**—**Check the **None** check box to disable rekey, check the **SSL** check box to specify SSL renegotiation during a rekey, or check the **New Tunnel** check box to establish a new tunnel during rekey.

Modes

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

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

Dead Peer Detection

Dead Peer Detection (DPD) ensures that the security appliance (gateway) or the client can quickly detect a condition where the peer is not responding, and the connection has failed.

Fields

- Gateway Side Detection**—**Uncheck the **Disable** check box to specify that DPD is performed by the security appliance (gateway). Enter the interval, from 30 to 3600 seconds, with which the security appliance performs DPD.

- Client Side Detection—Uncheck the **Disable** check box to specify that DPD is performed by the client. Enter the interval, from 30 to 3600 seconds, with which the client performs DPD.

Modes

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

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

Customization

Fields

- Portal Customization—Selects the customization to apply to the AnyConnect Client/SSL VPN portal page. The default is DfltCustomization.
- Manage—Opens the Configure GUI Customization objects dialog box, in which you can specify that you want to add, edit, delete, import, or export a customization object.
- Access Deny Message—Specifies a message to display to the end user when the connection is denied. Select Inherit to accept the message in the default group policy. The default message, if you deselect Inherit, is: “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
•	—	•	—	—

ACLs

This dialog box lets you configure ACLs for Clientless SSL VPN.

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 ACLs and their ACEs in priority order according to their position in the ACLs list 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
•	—	•	—	—

Add/Edit Internal Group Policy > IPsec Client

The Add or Edit Group Policy > IPsec dialog box lets you specify tunneling protocols, filters, connection settings, and servers for the group policy being added or modified.

Fields

- Re-Authentication on IKE Re-key—Enables or disables reauthentication when IKE re-key occurs, unless the Inherit check box is checked. The user has 30 seconds to enter credentials, and up to three attempts before the SA expires at approximately two minutes and the tunnel terminates.
- Allow entry of authentication credentials until SA expires—Allow users the time to reenter authentication credentials until the maximum lifetime of the configured SA.
- IP Compression—Enables or disables IP Compression, unless the Inherit check box is checked.
- 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.
- 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.
- Tunnel Group Lock—Enables locking the tunnel group you select from the list, unless the Inherit check box or the value None is selected.
- 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
•	—	•	—	—

Client Access Rules

The Client Access Rules table on this dialog box lets you view up to 25 client access rules. If you uncheck 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 column contains a comma-separated list of software or firmware images 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
•	—	•	—	—

Add/Edit Client Access Rule

The Add or Edit Client Access Rule dialog box 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 Dialog Box

The Add or Edit Group Policy dialog box, Client Configuration dialog box contains three tabs that let you configure general client parameters, Cisco client parameters, and Microsoft client parameters.

For information about the individual dialog boxes, see the following links:

- Add/Edit Internal Group Policy > Client Configuration > General Client Parameters, page 36-32
- Add/Edit Internal Group Policy > Client Configuration > Cisco Client Parameters, page 36-33
- Add or Edit Internal Group Policy > Advanced > IE Browser Proxy, page 36-34

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 > General Client Parameters

This dialog box configures client attributes that are common across both Cisco and Microsoft clients, including the banner text, default domain, split tunnel parameters, and address pools.

Note The AnyConnect VPN client and the SSL VPN client do not support split DNS.

Fields

- Inherit—(Multiple instances) Indicates that the corresponding setting takes its value from the default group policy. Unchecking 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 choose from the drop-down list.
- Manage—Opens the ACL Manager dialog box, in 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. Unchecking 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

The View/Config Banner dialog box lets you enter up to 500 characters of text in the text field to display 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 > Cisco Client Parameters

This dialog box 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 or Edit Internal Group Policy > Advanced > IE Browser Proxy

This dialog box configures attributes 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.
  - Select proxy server settings from the following—Enables the following check boxes for your selections: Auto detect proxy, Use proxy server settings given below, and Use proxy auto configuration (PAC) given below.
  - 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.
  - Use proxy auto configuration (PAC) given below—Specifies the use of the file specified in the Proxy Auto Configuration (PAC) field as the source for auto configuration attributes.
- Proxy Server Settings—Configures the proxy server parameters for Microsoft clients using Microsoft Internet Explorer.
  - Server Address and Port—Specifies the IP address or name and the port of an Microsoft Internet Explorer server that is applied for this client PC.
  - Bypass Proxy Server for Local Addresses—Configures Microsoft Internet Explorer browser proxy local-bypass settings for a client PC. Click **Yes** to enable local bypass or **No** to disable local bypass.

    - Exception List—Lists the server names and IP addresses that you want to exclude from proxy server access. Enter the list of addresses that you do not want to have accessed through a proxy server. This list corresponds to the Exceptions list in the Proxy Settings dialog box in Internet Explorer.

- PAC URL—Specifies the URL of the auto-configuration file. This file tells the browser where to look for proxy information. To use the proxy auto-configuration (PAC) feature, the remote user must use the Cisco AnyConnect VPN client.

  Many network environments define HTTP proxies that connect a web browser to a particular network resource. The HTTP traffic can reach the network resource only if the proxy is specified in the browser and the client routes the HTTP traffic to the proxy. SSLVPN tunnels complicate the definition of HTTP proxies because the proxy required when tunneled to an enterprise network can differ from that required when connected to the Internet via a broadband connection or when on a third-party network.

  In addition, companies with large networks might need to configure more than one proxy server and let users choose between them, based on transient conditions. By using .pac files, an administrator can author a single script file that determines which of numerous proxies to use for all client computers throughout the enterprise.

  The following are some examples of how you might use a PAC file:

    - Choosing a proxy at random from a list for load balancing.
    - Rotating proxies by time of day or day of the week to accommodate a server maintenance schedule.
    - Specifying a backup proxy server to use in case the primary proxy fails.
    - Specifying the nearest proxy for roaming users, based on the local subnet.

  You can use a text editor to create a proxy auto-configuration (.pac) file for your browser. A .pac file is a JavaScript file that contains logic that specifies one or more proxy servers to be used, depending on the contents of the URL. Use the PAC URL field to specify the URL from which to retrieve the .pac file. Then the browser uses the .pac file to determine the proxy settings. For details about .pac files, see the following Microsoft Knowledge Base article: http://www.microsoft.com/mind/0599/faq/faq0599.asp.

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

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. Choose 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

The Add or Edit Group Policy Client Firewall dialog box 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 dialog boxing 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 dialog box are active. If you want users in this group to be firewall-protected, select either the Firewall Required or Firewall Optional setting.

  If you choose 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 choose 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 dialog box, in 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

The Add or Edit Group Policy > Hardware Client dialog box lets you configure settings for the VPN 3002 hardware client for the group policy being added or modified. The Hardware Client dialog box 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 dialog box.
- 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 dialog box for the hardware client. To authenticate, click Connect/Login Status.

  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.

  Note that the idle timeout indicated in response to the show uauth command is always the idle timeout value of the user who authenticated the tunnel on the Cisco Easy VPN remote device.

- 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 C—Restricts the use of Network Extension Mode on the hardware client. Choose 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:

Add/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, in which you can configure a new server or URL and display name.
- Edit—Opens the Edit Server or URL dialog box, in 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

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.

Configuring AnyConnect (SSL) VPN Client Connections

The Cisco AnyConnect VPN client provides secure SSL connections to the security appliance for remote users. The client gives remote users the benefits of an SSL VPN client without the need for network administrators to install and configure clients on remote computers.

Without a previously-installed client, remote users enter the IP address in their browser of an interface configured to accept SSL VPN connections. Unless the security appliance is configured to redirect http:// requests to https://, users must enter the URL in the form https://<*address*>.

After entering the URL, the browser connects to that interface and displays the login screen. If the user satisfies the login and authentication, and the security appliance identifies the user as requiring the client, it downloads the client that matches the operating system of the remote computer. After downloading, the client installs and configures itself, establishes a secure SSL connection and either remains or uninstalls itself (depending on the security appliance configuration) when the connection terminates.

In the case of a previously installed client, when the user authenticates, the security appliance examines the revision of the client, and upgrades the client as necessary.

When the client negotiates an SSL VPN connection with the security appliance, it connects using Transport Layer Security (TLS), and optionally, Datagram Transport Layer Security (DTLS). DTLS avoids latency and bandwidth problems associated with some SSL connections and improves the performance of real-time applications that are sensitive to packet delays.

The AnyConnect client can be downloaded from the security appliance, or it can be installed manually on the remote PC by the system administrator. For more information about installing the client manually, see the *Cisco AnyConnect VPN Client Administrator Guide*.

The security appliance downloads the client based on the group policy or username attributes of the user establishing the connection. You can configure the security appliance to automatically download the client, or you can configure it to prompt the remote user about whether to download the client. In the latter case, if the user does not respond, you can configure the security appliance to either download the client after a timeout period or present the login page.

Fields

- Keep Installer on Client System—Enable to allow permanent client installation on the remote computer. Enabling disables the automatic uninstalling feature of the client. The client remains installed on the remote computer for subsequent connections, reducing the connection time for the remote user.
- Compression—Compression increases the communications performance between the security appliance and the client by reducing the size of the packets being transferred.
- Datagram TLS—DTLS avoids latency and bandwidth problems associated with some SSL connections and improves the performance of real-time applications that are sensitive to packet delays.
- Ignore Defragmentation (DF) Bit—By default, the security appliance discards SSL-encapsulated packets that exceed the SSL MTU. IPsec does not have an MTU, so an SSL session cannot, by default, encapsulate IPsec. If you want to support IPsec within an SSL session, enable this parameter to prevent the security appliance from discarding packets that exceed the SSL MTU. You must also enable the Ignore Routing and Filtering Rules parameter. An example use case is to let users establish an SSL VPN session with the security appliance and use that session to establish an IPsec VPN session with another enterprise. Without the initial SSL VPN session, network policies might prevent the establishment of an IPsec session.
- Ignore Routing and Filtering Rules—By default, the group policy pushed to the SSL client permits client enforcement of routing and filtering rules configured on the endpoint. These rules can prevent the transmission of SSL-encapsulated packets containing IPsec-encapsulated packets. For example, the client may have a rule that prevents the exchange of SSL-encapsulated packets that exceed the MTU size. If you want to support IPsec within an SSL session, enable this parameter to prevent the client from discarding packets that exceed the SSL MTU. You must also enable the Ignore Defragmentation (DF) Bit parameter.

- Keepalive Messages—Enter an number, from 15 to 600 seconds, in the Interval field to enable and adjust the interval of keepalive messages to ensure that an 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 interval also ensures that the client 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.
- MTU—Adjusts the MTU size for SSL connections. Enter a value in bytes, from 256 to 1410 bytes. By default, the MTU size is adjusted automatically based on the MTU of the interface that the connection uses, minus the IP/UDP/DTLS overhead.
- Client Profile to Download—a profile is a group of configuration parameters that the AnyConnect client uses to configure the connection entries that appear in the user interface, including the names and addresses of host computers.
- **Optional Client Module to Download**—To minimize download time, the AnyConnect client requests downloads (from the security appliance) only of modules that it needs for each feature that it supports. You must specify the names of modules that enable other features:
  - Enable the Start Before Logon (SBL) feature by checking **vpngina**. This enables the security appliance to download a graphical identification and authentication (GINA) for the AnyConnect client VPN connection.
  - Enable the Cisco Diagnostic AnyConnect Reporting Tool (DART) by checking **dart**. DART captures a snapshot of system logs and other diagnostic information and creates a .zip file on your desktop so you can conveniently send troubleshooting information to Cisco TAC. For this keyword to have any effect, you must have installed the DART package 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
•	—	•	—	—

Configuring SSL VPN Connections

Use the AnyConnect Connection Profiles pane and its child dialog boxes to specify SSL VPN connection attributes for client-based connections. These attributes apply to the Cisco AnyConnect VPN client and to the legacy SSL VPN client.

The initial client deployment requires end-user administrative rights. The Cisco AnyConnect VPN client supports the HTTPS/TCP (SSL) and Datagram Transport Layer Security (DTLS) tunneling options.

In the main pane, you can enable client access on the interfaces you select and you can select, add, edit, and delete connections (tunnel groups). You can also specify whether you want to allow a user to select a particular connection at login.

Fields

- Access Interfaces—Specify SSL VPN client access for each interface listed in the table:

  - Enable Cisco AnyConnect VPN Client or legacy SSL VPN Client access on the interfaces in the table below—Check this check box to enable the AnyConnect VPN client or the legacy SSL VPN client on the interfaces listed in the Access Interfaces table.
  - Interface—The interface to enable SSL VPN client connections.
  - Allow Access—Check Allow Access to enable access on the interfaces listed in this table.
  - Enable DTLS—Check Enable DTLS to enable Datagram Transport Layer Security (DTLS) on an interface. DTLS avoids latency and bandwidth problems associated with some SSL connections and improves the performance of real-time applications that are sensitive to packet delays.
  - Access Port—Specify the port for SSL VPN client connections. The default value is 443.
  - DTLS Port—Specify the port for DTLS connections. The default value is 443.
- Login Page Setting—Allow the user to select a connection profile, identified by its alias, on the login page. If you do not check this check box, the default connection profile is DefaultWebVPNGroup.
- Connection Profiles—Configure protocol-specific attributes for connections (tunnel groups).
  - Add/Edit—Click to Add or Edit a Connection Profile (tunnel group).
  - Name—The name of the Connection Profile.
  - Aliases—Other names by which the Connection Profile is known.
  - SSL VPN Client Protocol—Specifies whether SSL VPN client have access.
  - Group Policy—Shows the default group policy for this Connection Profile.
  - Allow user to select connection, identified by alias in the table above, at login page—Check to enable the display of Connection Profile (tunnel group) aliases on the Login page.

Modes

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

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

Setting the Basic Attributes for an SSL VPN Connection

To set the basic attributes for an SSL VPN connection, choose Add or Edit in the Connection Profiles section. The Add (or Edit) SSL VPN Connection > Basic dialog box opens.

Fields

Set the attributes in the Add SSL VPN Connection > Basic dialog box as follows:

- Name—For Add, specify the name of the connection profile you are adding. For Edit, this field is not editable.
- Aliases—(Optional) Enter one or more alternative names for the connection. You can spaces or punctuation to separate the names.

- Authentication—Choose one of the following methods to use to authenticate the connection and specify a AAA server group to use in authentication.
  - AAA, Certificate, or Both—Select the type of authentication to use: AAA, Certificate, or Both. If you choose either Certificate or Both, the user must provide a certificate in order to connect.
  - AAA Server Group—Choose a AAA server group from the drop-down list. The default setting is LOCAL, which specifies that the security appliance handles the authentication. Before making a selection, you can click **Manage** to open a dialog box over this dialog box to view or make changes to the security appliance configuration of AAA server groups.
  - Choosing something other than LOCAL makes available the Use LOCAL if Server Group Fails check box.
  - Use LOCAL if Server Group fails—Check to enable the use of the LOCAL database if the group specified by the Authentication Server Group attribute fails.
- Client Address Assignment—Select the DHCP servers, client address pools, and client IPv6 address pools to use.
  - DHCP Servers—Enter the name or IP address of a DHCP server to use.
  - Client Address Pools—Enter the pool name of an available, configured pool of IP addresses to use for client address assignment. Before making a selection, you can click **Select** to open a dialog box over this dialog box to view or make changes to the address pools.
- Default Group Policy—Select the group policy to use.
  - Group Policy—Select the VPN group policy that you want to assign as the default group policy for this connection. A VPN group policy is a collection of user-oriented attribute-value pairs that can be stored internally on the device or externally on a RADIUS server. The default value is DfltGrpPolicy. You can click **Manage** to open a dialog box over this one to make changes to the group policy configuration.
  - Enable SSL VPN Client Protocol—Check the check box to enable SSL VPN for this connection; uncheck to disable it.
- Find—Enter a GUI label or a CLI command to use as a search string, then click Next or Previous to begin the search.

Modes

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

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

Setting Advanced Attributes for a Connection Profile

The Advanced menu items and their dialog boxes let you configure the following characteristics for this connection:

- General attributes.
- Client Addressing attributes
- Authentication attributes.

- Authorization attributes.
- Accounting attributes.
- Name server attributes.
- Clientless SSL VPN attributes.

Note SSL VPN and secondary authentication attributes apply only to SSL VPN connection profiles.

Setting General Attributes for an AnyConnect SSL VPN Connection

Configure the General attributes to specify the password management parameters.

Fields

Set the Advanced General attributes as follows:

- Enable Password Management—Lets you configure parameters relevant to overriding an account-disabled indication from a AAA server and to notifying users about password expiration.

  The security appliance supports password management for the RADIUS and LDAP protocols. It supports the "password-expire-in-days" option only for LDAP. This parameter is valid for AAA servers that support such notification. The security appliance ignores this command if RADIUS or LDAP authentication has not been configured.

  You can configure password management for IPsec remote access and SSL VPN tunnel-groups.

  Note Some RADIUS servers that support MS-CHAP currently do not support MS-CHAPv2. This feature requires MS-CHAPv2, so please check with your vendor.

  The security appliance, releases 7.1 and later, generally supports password management for the following connection types when authenticating with LDAP or with any RADIUS configuration that supports MS-CHAPv2:

  - AnyConnect VPN client
  - IPsec VPN client
  - Clientless SSL VPN

  Password management is *not* supported for any of these connection types for Kerberos/Active Directory (Windows password) or NT 4.0 Domain. The RADIUS server (for example, Cisco ACS) could proxy the authentication request to another authentication server. However, from the security appliance perspective, it is talking only to a RADIUS server.

  Note For LDAP, the method to change a password is proprietary for the different LDAP servers on the market. Currently, the security appliance implements the proprietary password management logic only for Microsoft Active Directory and Sun LDAP servers.

  Native LDAP requires an SSL connection. You must enable LDAP over SSL before attempting to do password management for LDAP. By default, LDAP uses port 636.

  Note Allowing override account-disabled is a potential security risk.

- Notify user __ days prior to password expiration—Specifies that ASDM must notify the user at login a specific number of days before the password expires. The default is to notify the user 14 days prior to password expiration and every day thereafter until the user changes the password. The range is 1 through 180 days.
- Notify user on the day password expires—Notifies the user only on the day that the password expires.

  In either case, and, if the password expires without being changed, the security appliance offers the user the opportunity to change the password. If the current password has not expired, the user can still log in using that password.

Note This does not change the number of days before the password expires, but rather, it enables the notification. If you select this option, you must also specify the number of days.

- Override account-disabled indication from AAA server—Overrides an account-disabled indication from a AAA server.

- Find—Enter a GUI label or a CLI command to use as a search string, then click Next or Previous to begin the search.

Modes

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

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

Setting Client Addressing Attributes for an AnyConnect SSL VPN Connection

The Client Addressing attributes let you configure interface-specific address pools that your connection can use. Click Add to add a new address pool or Edit to modify an existing pool. The Select Address Pools dialog box opens, showing a table listing the pool name, starting and ending address (or number of addresses), and subnet mask/prefix length of any existing pools. For a complete description of Client Addressing see Configuring Client Addressing, page 36-76.

Configuring Authentication Attributes for a Connection Profile

- Interface-specific Authentication Server Groups—Manages the assignment of authentication server groups to specific interfaces.
  - Add or Edit—Opens the Assign Authentication Server Group to Interface dialog box, in which you can specify the interface and server group, and specify whether to allow fallback to the LOCAL database if the selected server group fails. The Manage button on this dialog box opens the Configure AAA Server Groups dialog box. Your selections appear in the Interface/Server Group table.
  - Delete—Removes the selected server group from the table. There is no confirmation or undo.

- Username Mapping from Certificate—Lets you specify the methods and fields in a digital certificate from which to extract the username.
  - Pre-fill Username from Certificate—Extracts the username from the specified certificate field and uses it for username/password authentication and authorization, according to the options that follow in this panel.
  - Hide username from end user—Specifies to not display the extracted username to the end user.
  - Use script to select username—Specify the name of a script to use to select a username from a digital certificate. The default is --None--.
  - Add or Edit—Opens the Add or Edit Script Content dialog box, in which you can define a script to use in mapping the username from the certificate.
  - Delete—Deletes the selected script. There is no confirmation or undo.
  - Use the entire DN as the username—Specifies that you want to use the entire Distinguished Name field of the certificate as the username.
  - Specify the certificate fields to be used as the username—Specifies one or more fields to combine into the username.

    Possible values for primary and secondary attributes include the following:

Attribute	Definition
C	Country: the two-letter country abbreviation. These codes conform to ISO 3166 country abbreviations.
CN	Common Name: the name of a person, system, or other entity. Not available a s a secondary attribute.
DNQ	Domain Name Qualifier.
EA	E-mail address.
GENQ	Generational Qualifier.
GN	Given Name.
I	Initials.
L	Locality: the city or town where the organization is located.
N	Name.
O	Organization: the name of the company, institution, agency, association or other entity.
OU	Organizational Unit: the subgroup within the organization (O).
SER	Serial Number.
SN	Surname.
SP	State/Province: the state or province where the organization is located
T	Title.
UID	User Identifier.
UPN	User Principal Name.

  - Primary Field—Selects the first field to use from the certificate for the username. If this value is found, the secondary field is ignored.
  - Secondary Field—Selects the field to us if the primary field is not found.

- Find—Enter a GUI label or a CLI command to use as a search string, then click Next or Previous to begin the search.

Modes

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

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

Configuring Secondary Authentication Attributes for an SSL VPN Connection Profile

The Secondary Authentication dialog box lets you configure secondary or “double” authentication for this connection profile. With double authentication enabled, the end user must present two sets of valid authentication credentials in order to log on. You can use secondary authentication in conjunction with pre-filling the username from a certificate. The fields in this dialog box are similar to those you configure for primary authentication, but these fields relate only to secondary authentication.

When double authentication is enabled, these attributes select one or more fields in a certificate to use as the username. Configuring the secondary username from certificate attribute forces the security appliance to use the specified certificate field as the second username for the second username/password authentication.

Note If you also specify the secondary authentication server group, along with the secondary username from certificate, only the primary username is used for authentication.

Fields

- Secondary Authorization Server Group—Specifies an authorization server group from which to extract secondary credentials.
  - Server Group—Select an authorization server group to use as the secondary server AAA group. The default is none. The secondary server group cannot be an SDI server group.
  - Manage—Opens the Configure AAA Server Groups dialog box.
  - Use LOCAL if Server Group fails—Specifies to fall back to the LOCAL database if the specified server group fails.
- Use primary username—Specifies that the login dialog must request only one username.
  - Attributes Server—Select whether this is the primary or secondary attributes server.

    Note If you also specify an authorization server for this connection profile, the authorization server settings take precedence—the security appliance ignores this secondary authentication server.

  - Session username Server—Select whether this is the primary or secondary session username server.

- Interface-specific Authorization Server Groups—Manages the assignment of authorization server groups to specific interfaces.
  - Add or Edit—Opens the Assign Authentication Server Group to Interface dialog box, in which you can specify the interface and server group, and specify whether to allow fallback to the LOCAL database if the selected server group fails. The Manage button on this dialog box opens the Configure AAA Server Groups dialog box. Your selections appear in the Interface/Server Group table.
  - Delete—Removes the selected server group from the table. There is no confirmation or undo.
- Specify the certificate fields to be used as the username—Specifies one or more fields to match as the username. To use this username in the pre-fill username from certificate feature for the secondary username/password authentication or authorization, you must also configure the pre-fill-username and secondary-pre-fill-username.

  Possible values for primary and secondary attributes include the following:

Attribute	Definition
C	Country: the two-letter country abbreviation. These codes conform to ISO 3166 country abbreviations.
CN	Common Name: the name of a person, system, or other entity. Not available a s a secondary attribute.
DNQ	Domain Name Qualifier.
EA	E-mail address.
GENQ	Generational Qualifier.
GN	Given Name.
I	Initials.
L	Locality: the city or town where the organization is located.
N	Name.
O	Organization: the name of the company, institution, agency, association or other entity.
OU	Organizational Unit: the subgroup within the organization (O).
SER	Serial Number.
SN	Surname.
SP	State/Province: the state or province where the organization is located
T	Title.
UID	User Identifier.
UPN	User Principal Name.

  - Primary Field—Selects the first field to use from the certificate for the username. If this value is found, the secondary field is ignored.
  - Secondary Field—Selects the field to us if the primary field is not found.

Modes

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

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

Configuring Authorization Attributes for an SSL VPN Connection Profile

The Authorization dialog box lets you view, add, edit, or delete interface-specific authorization server groups. Each row of the table on this dialog box shows the status of one interface-specific server group: the interface name, its associated server group, and whether fallback to the local database is enabled if the selected server group fails.

Fields

- Authorization Server Group—Specifies an authorization server group from which to draw authorization parameters.
  - Server Group—Selects an authorization server group to use. The default is none.
  - Manage—Opens the Configure AAA Server Groups dialog box.
  - Users must exist in the authorization database to connect—Select this check box to require that users meet this criterion.
- Interface-specific Authorization Server Groups—Manages the assignment of authorization server groups to specific interfaces.
  - Add or Edit—Opens the Assign Authentication Server Group to Interface dialog box, in which you can specify the interface and server group, and specify whether to allow fallback to the LOCAL database if the selected server group fails. The Manage button on this dialog box opens the Configure AAA Server Groups dialog box. Your selections appear in the Interface/Server Group table.
  - Delete—Removes the selected server group from the table. There is no confirmation or undo.
- Username Mapping from Certificate—Specify the fields in a digital certificate from which to extract the username.
  - Use script to select username—Specifies the name of a script to use to select a username from a digital certificate. The default is --None--.
  - Add or Edit—Opens the Add or Edit Script Content dialog box, in which you can define a script to use in mapping the username from the certificate.
  - Delete—Deletes the selected script. There is no confirmation or undo.
  - Use the entire DN as the username—Specifies that you want to use the entire Distinguished Name field of the certificate as the username.
  - Specify the certificate fields to be used as the username—Specifies one or more fields to combine into the username.
  - Primary Field—Selects the first field to use in the certificate for the username. If this value is found, the secondary field is ignored.
  - Secondary Field—Selects the field to use if the primary field is not found.
- Find—Enter a GUI label or a CLI command to use as a search string, then click Next or Previous to begin the search.

Modes

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

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

Adding or Editing Content to a Script for Certificate Pre-Fill-Username

The Add or Edit Script Content dialog box lets you create an authentication or authorization script.

Note Both AnyConnect client and clientless WebVPN display "Unknown" in the username field when pre-fill-username from certificate using a script cannot find the username in the client certificate.

Fields

- Script Name—Specify the name of the script. The script name must be the same in both authorization and authentication.You define the script here, and CLI uses the same script to perform this function.
- Select script parameters—Specify the attributes and content of the script.
- Value for Username—Select an attribute from the drop-down list of standard DN attributes to use as the username (Subject DN).
- No Filtering—Specify that you want to use the entire specified DN name.
- Filter by substring— Specify the Starting Index (the position in the string of the first character to match) and Ending Index (number of characters to search). If you choose this option, the starting index cannot be blank. If you leave the ending index blank, it defaults to -1, indicating that the entire string is searched for a match.

  For example, suppose you selected the DN attribute Common Name (CN), which contains a value of host/user. Table 36-1 shows some possible ways you might filter this value using the substring option to achieve various return values. The Return Value is what is actually pre-filled as the username.

Table 36-1 Filtering by Substring

Starting Index	Ending Index	Return Value
1	5	host/
6	10	user
6	-1	user

Using a negative index, as in the third row of this table, specifies to count from the end of the string backwards to the end of the substring, in this case, the "r" of "user".

When using filtering by substrings, you should know the length of the substring that you are seeking. From the following examples, use either the regular expression matching or the custom script in Lua format:

- Example 1: Regular Expression Matching—Enter a regular expression to apply to the search in the Regular Expression field. Standard regular expression operators apply. For example, suppose you want to use a regular expression to filter everything up to the @ symbol of the "Email Address (EA)" DN value. The regular expression ^[^@]* would be one way to do this. In this example, if the DN value contained a value of user1234@company.com, the return value after the regular expression would be user1234.
- Example 2: Use custom script in Lua format—Specify a custom script written in the Lua programming language to parse the search fields. Selecting this option makes available a field in which you can enter your custom Lua script; for example, the script:

  ```
  return cert.subject.cn..'/'..cert.subject.l
  ```

  combines two DN fields, username (cn) and locality (l), to use as a single username and inserts the slash (/) character between the two fields.

  Table 36-2 lists the attribute names and descriptions that you can use in a Lua script.

Note Lua is case-sensitive.

Table 36-2 Attribute Names and Descriptions

Attribute Name	Description
cert.subject.c	Country
cert.subject.cn	Common Name
cert.subject.dnq	DN qualifier
cert.subject.ea	Email Address
cert.subject.genq	Generational qualified
cert.subject.gn	Given Name
cert.subject.i	Initials
cert.subject.l	Locality
cert.subject.n	Name
cert.subject.o	Organization
cert.subject.ou	Organization Unit
cert.subject.ser	Subject Serial Number
cert.subject.sn	Surname
cert.subject.sp	State/Province
cert.subject.t	Title
cert.subject.uid	User ID
cert.issuer.c	Country
cert.issuer.cn	Common Name
cert.issuer.dnq	DN qualifier
cert.issuer.ea	Email Address
cert.issuer.genq	Generational qualified
cert.issuer.gn	Given Name

Table 36-2 Attribute Names and Descriptions

cert.issuer.i	Initials
cert.issuer.l	Locality
cert.issuer.n	Name
cert.issuer.o	Organization
cert.issuer.ou	Organization Unit
cert.issuer.ser	Issuer Serial Number
cert.issuer.sn	Surname
cert.issuer.sp	State/Province
cert.issuer.t	Title
cert.issuer.uid	User ID
cert.serialnumber	Certificate Serial Number
cert.subjectaltname.upn	User Principal Name

If an error occurs while activating a tunnel group script, causing the script not to activate, the administrator’s console displays an error message.

Modes

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

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

Configuring Clientless SSL VPN Connections

Use the Clientless SSL VPN Access Connections dialog box to configure clientless SSL VPN access parameters. This dialog box also records the configuration choices you make in its child dialog boxes.

Fields

- Access Interfaces—Lets you select from a table the interfaces on which to enable access. The fields in this table include the interface name and check boxes enabling you whether to allow access and require a certificate for authentication.
- Access Port—Specifies the access port for the connection. The default value is 443.
- Connections—Provides a connection table that shows the records that determine the connection policy for this connection (tunnel group). Each record identifies a default group policy for the connection and contains protocol-specific connection parameters.
  - Add—Opens the Add Clientless SSL VPN dialog box for the selected connection.
  - Edit—Opens the Edit Clientless SSL VPN dialog box for the selected connection.
  - Delete—Removes the selected connection from the table. There is no confirmation or undo.

  - Allow user to select connection, identified by alias in the table above, at login page—Specifies that the user login page presents the user with a drop-down menu from which the user can select a particular tunnel group with which to connect.

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 Clientless SSL VPN Connections

The Add or Edit SSL VPN dialog box consists of Basic and Advanced sections, accessible through the expandable menu on the left of the box.

Add or Edit Clientless SSL VPN Connections > Basic

The Basic dialog box lets you configure essential characteristics for this connection.

Fields

- Name—Specifies the name of the connection. For the edit function, this field is read-only.
- Aliases—(Optional) Specifies one or more alternate names for this connection. The aliases appear on the login page if you configure that option on the Clientless SSL VPN Access Connections dialog box.
- Authentication—Specifies the authentication parameters.
  - Method—Specifies whether to use AAA authentication, certificate authentication, or both methods for this connection. The default is AAA authentication.
  - AAA server Group—Selects the AAA server group to use for authenticating this connection. The default is LOCAL.
  - Manage—Opens the Configure AAA Server Groups dialog box.
- DNS Server Group—Selects the server to use as the DNS server group for this connection. The default is DefaultDNS.
- Default Group Policy—Specifies the default group policy parameters to use for this connection.
  - Group Policy—Selects the default group policy to use for this connection. The default is DfltGrpPolicy.
  - Clientless SSL VPN Protocol—Enables or disables the Clientless SSL VPN protocol for this connection.

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 Clientless SSL VPN Connections > Advanced

The Advanced menu items and their dialog boxes let you configure the following characteristics for this connection:

- General attributes.
- Authentication attributes.
- Authorization attributes.
- Accounting attributes.
- Name server attributes.
- Clientless SSL VPN attributes.

Add or Edit Clientless SSL VPN Connections > Advanced > General

Use this dialog box to specify whether to strip the realm and group from the username before passing them to the AAA server, and to specify password management options.

Fields

- Password Management—Lets you configure parameters relevant to overriding an account-disabled indication from a AAA server and to notifying users about password expiration.
  - Enable notification password management—Checking this check box makes the following two parameters available. You can select either to notify the user at login a specific number of days before the password expires or to notify the user only on the day that the password expires. The default is to notify the user 14 days prior to password expiration and every day thereafter until the user changes the password. The range is 1 through 180 days.

    Note This does not change the number of days before the password expires, but rather, it enables the notification. If you select this option, you must also specify the number of days.

    In either case, and, if the password expires without being changed, the security appliance 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.
  - 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.

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 Clientless or SSL VPN Client Connection Profile or IPsec Connection Profiles> Advanced > Authentication

The Authentication dialog box lets you view, add, edit, or delete interface-specific authentication server groups. Each row of the table on this dialog box shows the status of one interface-specific server group: the interface name, its associated server group, and whether fallback to the local database is enabled if the selected server group fails.

Fields

- Interface-specific Authorization Server Groups—Manages the assignment of authorization server groups to specific interfaces.
  - Add or Edit—Opens the Assign Authentication Server Group to Interface dialog box, in which you can specify the interface and server group, and specify whether to allow fallback to the LOCAL database if the selected server group fails. The Manage button on this dialog box opens the Configure AAA Server Groups dialog box. Your selections appear in the Interface/Server Group table.
  - Delete—Removes the selected server group 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
•	—	•	—	—

Assign Authentication Server Group to Interface

This dialog box lets you associate an interface with a AAA server group. The results appear in the table on the Authentication dialog box.

Fields

- Interface—Selects an interface, DMZ, Outside, or Inside. The default is DMZ.
- Server Group—Selects a server group to assign to the selected interface. The default is LOCAL.
- Manage—Opens the Configure AAA Server Groups dialog box.
- Fallback—Enables or disables fallback to LOCAL if the selected server group fails.

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 SSL VPN Connections > Advanced > Authorization

This dialog box lets you configure the default authorization server group, interface-specific authorization server groups, and user name mapping attributes. The attributes are the same for SSL VPN and Clientless SSL VPN connections.

Fields

- Default Authorization Server Group—Configures default authorization server group attributes.
  - Server Group—Selects the authorization server group to use for this connection. The default is --None--.
  - Manage—Opens the Configure AAA Server Groups dialog box.
  - Users must exist in the authorization database to connect—Enables or disables this requirement.
- Interface-specific Authorization Server Groups
  - Table—Lists each configured interface and the server group with which it is associated.
  - Add or Edit—Opens the Assign Authorization Server Group to Interface dialog box.
  - Delete—Removes the selected row from the table.
- User Name Mapping—Specifies user name mapping attributes.
- Username Mapping from Certificate—Lets you specify the fields in a digital certificate from which to extract the username.
  - Pre-fill Username from Certificate —Enables the use of a username extracted from the specified certificate field as the username for username/password authentication and authorization, using the options that follow in this dialog box.
  - Hide username from end user—Specifies not to display the extracted username to the end user.
  - Use script to select username—Specify the name of a script to use to select a username from a digital certificate. There is no default.
  - Add or Edit—Opens the Add or Edit Script Content dialog box, in which you can define a script to use in mapping the username from the certificate.
  - Delete—Deletes the selected script. There is no confirmation or undo.

- Use the entire DN as the username—Enables or disables the requirement to use the entire DN as the username.
- Specify individual DN fields as the username. You can select both the primary DN field, for which the default is CN (Common Name) and the secondary DN field, for which the default is OU (Organization Unit).
- Primary Field—Selects the first field to use in the username.
- Secondary Field—Selects the second field to use in the username.

Modes

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

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

Assign Authorization Server Group to Interface

This dialog box lets you associate an interface with a AAA server group. The results appear in the table on the Authorization dialog box.

Fields

- Interface—Selects an interface, DMZ, Outside, or Inside. The default is DMZ.
- Server Group—Selects a server group to assign to the selected interface. The default is LOCAL.
- Manage—Opens the Configure AAA Server Groups 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 or Edit SSL VPN Connections > Advanced > SSL VPN

This dialog box lets you configure attributes that affect what the remote user sees upon login.

Fields

- Portal Page Customization—Configures the look and feel of the user login page by specifying which preconfigured customization attributes to apply. The default is DfltCustomization.
- Enable the display of Radius Reject-Message on the login screen—Select this check box to display the RADIUS-reject message on the login dialog box when authentication is rejected.

- Enable the display of SecurId message on the login screen—Select this check box to display SecurID messages on the login dialog box.
- Manage—Opens the Configure GUI Customization Objects dialog box.
- Connection Aliases—Lists in a table the existing connection aliases and their status and lets you add or delete items in that table. A connection alias appears on the user login page if the connection is configured to allow users to select a particular connection (tunnel group) at login. The rows in this table are editable in place, so there is no Edit button. Clicking the "i" icon above the table opens a tooltip for the edit function.
  - Add—Opens the Add Connection Alias dialog box, on which you can add and enable a connection alias.
  - Delete—Removes the selected row from the connection alias table. There is no confirmation or undo.
  - To edit an alias listed in the table, double-click the line.
- Group URLs—Lists in a table the existing group URLs and their status and lets you add or delete items in that table. A group URL appears on the user login page if the connection is configured to allow users to select a particular group at login. The rows in this table are editable in place, so there is no Edit button. Clicking the "i" icon above the table opens a tooltip for the edit function.
  - Add—Opens the Add Group URL dialog box, on which you can add and enable a group URL.
  - Delete—Removes the selected row from the connection alias table. There is no confirmation or undo.
  - To edit a URL listed in the table, double-click the line.
- Do not run Cisco Secure Desktop (CSD) on client machine when using group URLs defined above to access the ASA. (If a client connects using a connection alias, this setting is ignored.)—Check if you want to exempt users from running CSD who use a URL that matches an entry in the Group URLs table. Be aware that doing so stops the security appliance from receiving endpoint criteria from these users, so you might have to change the DAP configuration to provide them with VPN 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 or Edit Clientless SSL VPN Connections > Advanced > Clientless SSL VPN

This dialog box lets you configure attributes that affect what the remote user sees upon login.

Fields

- Portal Page Customization—Configures the look and feel of the user login page by specifying which preconfigured customization attributes to apply. The default is DfltCustomization.
- Enable the display of Radius Reject-Message on the login screen—Select this check box to display the RADIUS-reject message on the login dialog box when authentication is rejected.

- Enable the display of SecurId message on the login screen—Select this check box to display SecurID messages on the login dialog box.
- Manage—Opens the Configure GUI Customization Objects dialog box.
- Connection Aliases—Lists in a table the existing connection aliases and their status and lets you add or delete items in that table. A connection alias appears on the user login page if the connection is configured to allow users to select a particular connection (tunnel group) at login.
  - Add—Opens the Add Connection Alias dialog box, on which you can add and enable a connection alias.
  - Delete—Removes the selected row from the connection alias table. There is no confirmation or undo.
- Group URLs—Lists in a table the existing group URLs and their status and lets you add or delete items in that table. A group URL appears on the user login page if the connection is configured to allow users to select a particular group at login.
  - Add—Opens the Add Group URL dialog box, on which you can add and enable a group URL.
  - Delete—Removes the selected row from the connection alias table. There is no confirmation or undo.
- Do not run Cisco Secure Desktop (CSD) on client machine when using group URLs defined above to access the ASA. (If a client connects using a connection alias, this setting is ignored.)—Check if you want to exempt users from running CSD who use a URL that matches an entry in the Group URLs table. Be aware that doing so stops the security appliance from receiving endpoint criteria from these users, so you might have to change the DAP configuration to provide them with VPN 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 or Edit Clientless SSL VPN Connections > Advanced > NetBIOS Servers

The table on this dialog box shows the attributes of the already-configured NetBIOS servers. The Add or Edit Tunnel Group dialog box for Clientless SSL VPN access, NetBIOS dialog box, lets you configure the NetBIOS attributes for the tunnel group. Clientless SSL VPN 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. Clientless SSL VPN 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
•	—	•	—	—

Configure DNS Server Groups

This dialog box displays the configured DNS servers in a table, including the server group name, servers, timeout in seconds, number of retries allowed, and domain name. You can add, edit, or delete DNS server groups on this dialog box.

Fields

- Add or Edit—Opens the Add or Edit DNS Server Group dialog box.
- Delete—Removes the selected row from the table. There is no confirmation or undo.
- DNS Server Group—Selects the server to use as the DNS server group for this connection. The default is DefaultDNS.
- Manage—Opens the Configure DNS Server Groups 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 or Edit Clientless SSL VPN Connections > Advanced > Clientless SSL VPN

This dialog box lets you specify portal-related attributes for Clientless SSL VPN connections.

Fields

- Portal Page Customization—Selects the customization to apply to the user interface.
- Manage—Opens the Configure GUI Customization Objects 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
•	—	•	—	—

IPsec Remote Access Connection Profiles

The parameters in the IPsec Connection Profiles dialog box let you configure IPsec remote access connections. Most of the parameters in this section were formerly configured under tunnel groups. An IPsec connection represents a connection-specific record for IPsec and Clientless SSL VPN connections.

The IPsec group uses the IPsec connection parameters to create a tunnel. An IPsec connection can be either remote-access or Site-to-Site. 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 connection parameters take precedence over parameters set for users and groups.

The Clientless SSL VPN tunnel-group parameters are the parameters of the Clientless SSL VPN group that you want to apply to this IPsec connection. You configure Clientless SSL VPN access on the Configuration > Clientless SSL VPN dialog box.

Fields

- Access Interfaces—Selects the interfaces to enable for IPsec access. The default is that no access is selected.
- Connections—Shows in tabular format the configured parameters for existing IPsec connections. The Connections table contains records that determine connection policies. A record identifies a default group policy for the connection and contains protocol-specific connection parameters. The table contains the following columns:

- Name—Specifies the name or IP address of the IPsec connection.
- ID Certificate—Specifies the name of the ID certificate, if available.
- IPsec Protocol—Indicates whether the IPsec protocol is enabled. You enable this protocol on the Add or Edit IPsec Remote Access Connection, Basic dialog box.
- L2TP/IPsec Protocol—Indicates whether the L2TP/IPsec protocol is enabled. You enable this protocol on the Add or Edit IPsec Remote Access Connection, Basic dialog box.
- Group Policy—Indicates the name of the group policy for this IPsec connection.

- Add or Edit—Opens the Add or Edit IPsec Remote Access Connection Profile dialog box.
- Delete—Removes the selected server group 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 or Edit an IPsec Remote Access Connection Profile

The Add or Edit IPsec Remote Access Connection Profile dialog box has a navigation pane that lets you select basic or advanced elements to configure.

Add or Edit IPsec Remote Access Connection Profile Basic

The Add or Edit IPsec Remote Access Connection Profile Basic dialog box lets you configure common attributes for IPsec connections.

Fields

- Name—Identifies the name of the connection.
- IKE Peer Authentication—Configures IKE peers.
  - Pre-shared key—Specifies the value of the pre-shared key for the connection. The maximum length of a pre-shared key is 128 characters.
  - Identity Certificate—Selects the name of an identity certificate, if any identity certificates are configured and enrolled.
  - Manage—Opens the Manage Identity Certificates dialog box, on which you can add, edit, delete, export, and show details for a selected certificate.
- User Authentication—Specifies information about the servers used for user authentication. You can configure more authentication information in the Advanced section.
  - Server Group—Selects the server group to use for user authentication. the default is LOCAL. If you select something other than LOCAL, the Fallback check box becomes available.
  - Manage—Opens the Configure AAA Server Groups dialog box.

  - Fallback—Specifies whether to use LOCAL for user authentication if the specified server group fails.
- Client Address Assignment—Specifies attributes relevant to assigning client attributes.
  - DHCP Servers—Specifies the IP address of a DHCP server to use. You can add up to 10 servers, separated by spaces.
  - DHCP Link—Specifies the sub-option 5, the Link Selection Sub-option for the Relay Information Option 82, defined by RFC 3527. This should only be used with servers that support this RFC.
  - DHCP Subnet—Specifies the IPv4 Subnet Selection Option, defined by RFC 3011. This should only be used with servers that support this RFC.
  - Client Address Pools—Specifies up to 6 predefined address pools. To define an address pool, go to Configuration > Remote Access VPN > Network Client Access > Address Assignment > Address Pools.
  - Select—Opens the Select Address Pools dialog box.
- Default Group Policy—Specifies attributes relevant to the default group policy.
  - Group Policy—Selects the default group policy to use for this connection. The default is DfltGrpPolicy.
  - Manage—Opens the Configure Group Policies dialog box, from which you can add, edit, or delete group policies.
  - Client Protocols—Selects the protocol or protocols to use for this connection. By default, both IPsec and L2TP over IPsec are selected.

Modes

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

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

Mapping Certificates to IPsec or SSL VPN Connection Profiles

When the security appliance receives an IPsec connection request with client certificate authentication, it assigns a connection profile to the connection according to policies you configure. That policy can be to use rules you configure, use the certificate OU field, use the IKE identity (i.e. hostname, IP address, key ID), the peer IP address, or a default connection profile. For SSL connections, the security appliance only uses the rules you configure.

For IPsec or SSL connections using rules, the security appliance evaluates the attributes of the certificate against the rules until it finds a match. When it finds a match, it assigns the connection profile associated with the matched rule to the connection. If it fails to find a match, it assigns the default connection profile (DefaultRAGroup for IPsec and DefaultWEBVPNGroup for SSL VPN) to the connection and lets the user choose the connection profile from a drop-down menu displayed on the portal page (if it is enabled). The outcome of the connection attempt once in this connection profile depends on whether or not the certificate is valid and the authentication settings of the connection profile.

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.

First configure the policy for matching a certificate to a connection profile at Configuration > Remote Access VPN> Network (Client) Access > Advanced > IPSec > Certificate to Connection Profile Maps. If you choose to use rules you configure, go to Rules to specify the rules. The following procedures shows how you create the certificate-based criteria for each IPsec and SSL VPN connection profile:

Step 1 Use the table at the top (Certificate to Connection Profile Maps) to do one of the following:

- Create a list name, called a "map," specify the priority of the list, and assign the list to a connection profile.

  ASDM highlights the list after you add it to the table.

- Confirm that a list is assigned to the connection profile for which you want to add certificate-based rules.

  ASDM highlights the list after you add it to the table and displays any associated list entries in the table at the bottom of the pane.

Step 2 Use the table at the bottom (Mapping Criteria) to view, add, change or delete entries to the selected list.

Each entry in the list consists of one certificate-based rule. All of the rules in the mapping criteria list need to match the contents of the certificate for the security appliance to choose the associated map index. To assign a connection if one criterion or another matches, create one list for each matching criterion.

To understand the fields, see the following sections:

- Setting a Certificate Matching Policy
- Add/Edit Certificate Matching Rule
- Add/Edit Certificate Matching Rule Criterion

Setting a Certificate Matching Policy

For IPsec connections, 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.

Add/Edit Certificate Matching Rule

Use the **Add/Edit Certificate Matching Rule** dialog box to assign the name of a list (map) to a connection profile.

Fields

- **Map**—Choose one of the following:
  - **Existing**—Select the name of the map to include the rule.
  - **New**—Enter a new map name for a rule.
- **Rule Priority**—Type a decimal to specify the sequence with which the security appliance evaluates the map when it receives a connection request. For the first rule defined, the default priority is 10. The security appliance evaluates each connection against the map with the lowest priority number first.
- **Mapped to Connection Profile**—Select the connection profile, formerly called a "tunnel group," to map to this rule.

  If you do not assign a rule criterion to the map, as described in the next section, the security appliance ignores the map 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 Certificate Matching Rule Criterion

Use the **Add/Edit Certificate Matching Rule Criterion** dialog box to configure a certificate matching rule criterion for the selected connection profile.

Fields

- **Rule Priority**—(Display only). Sequence with which the security appliance evaluates the map when it receives a connection request. The security appliance evaluates each connection against the map with the lowest priority number first.
- **Mapped to Group**—(Display only). Connection profile to which the rule is assigned.
- **Field**—Select the part of the certificate to be evaluated from the drop-down list.
  - **Subject**—The person or system that uses the certificate. For a CA root certificate, the Subject and Issuer are the same.
  - **Alternative Subject**—The subject alternative names extension allows additional identities to be bound to the subject of the certificate.
  - **Issuer**—The CA or other entity (jurisdiction) that issued the certificate.
  - **Extended Key Usage**—An extension of the client certificate that provides further criteria that you can choose to match.

- **Component**—(Applies only if Subject of Issuer is selected.) 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.
Unstructured Name (UNAME)	The unstructuredName attribute type specifies the name or names of a subject as an unstructured ASCII string.
IP Address (IP)	IP address field.

- **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**—Enter up to 255 characters to specify the object of the operator. For Extended Key Usage, select one of the pre-defined values in the drop-down list, or you can enter OIDs for other extensions. The pre-defined values include the following:

Selection	Key Usage Purpose	OID String
clientauth	Client Authentication	1.3.6.1.5.5.7.3.2
codesigning	Code Signing	1.3.6.1.5.5.7.3.3
emailprotection	Secure Email Protection	1.3.6.1.5.5.7.3.4

Selection	Key Usage Purpose	OID String
clientauth	Client Authentication	1.3.6.1.5.5.7.3.2
ocspsigning	OCSP Signing	1.3.6.1.5.5.7.3.9
serverauth	Server Authentication	1.3.6.1.5.5.7.3.1
timestamping	Time Stamping	1.3.6.1.5.5.7.3.8

Modes

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

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

Configure Site-to-Site Tunnel Groups

The Tunnel Groups dialog box shows the attributes of the currently configured Site-to-Site tunnel groups, lets you select the delimiter to use when parsing tunnel group names, and lets you adds, modify, or delete tunnel groups.

Fields

- Add—Opens the Add IPsec Site-to-Site Tunnel Group dialog box.
- Edit—Opens the Edit IPsec Site-to-Site Tunnel Group dialog box.
- Delete—Removes the selected tunnel group. There is no confirmation or undo.
- Table of Tunnel Groups—Lists the tunnel group name, CA Certificate, IPsec protocol status (enabled or disabled), and group policy applied for each configured tunnel group.
- Group Delimiter—Selects the delimiter character to use parsing tunnel group names from the usernames that are received when tunnels are being negotiated.

Add/Edit Site-to-Site Connection

The Add or Edit IPsec Site-to-Site Connection dialog box lets you create or modify an IPsec Site-to-Site connection. These dialog boxes let you specify the peer IP address, specify a connection name, select an interface, specify IKE peer and user authentication parameters, specify protected networks, and specify encryption algorithms.

Fields

- Peer IP Address—Lets you specify an IP address and whether that address is static.
- Connection Name—Specifies the name assigned to this tunnel group. For the Edit function, this field is display-only. You can specify that the connection name is the same as the IP address specified in the Peer IP Address field.
- Interface—Selects the interface to use for this connection.

- IKE Authentication—Specifies the pre-shared key and ID certificate to use when authenticating an IKE peer.
  - Pre-shared Key—Specify the value of the pre-shared key for the tunnel group. The maximum length of the pre-shared key is 128 characters.
  - Identity Certificate—Specifies the name of the identity certificate, if available, to use for authentication.
  - Manage—Opens the Manage CA Certificates dialog box, on which you can see the certificates that are already configured, add new certificates, show details for a certificate, and edit or delete a certificate.
- Protected Networks—Selects or specifies the local and remote network protected for this connection.
  - Local Network—Specifies the IP address of the local network.
  - ...—Opens the Browse Local Network dialog box, in which you can select a local network.
  - Remote Network—Specifies the IP address of the remote network.
  - ...—Opens the Browse Remote Network dialog box, in which you can select a remote network.
- Encryption Algorithm—Specifies the encryption algorithms to use in the IKE and IPsec proposals.
  - IKE Proposal—Specifies one or more encryption algorithms to use for the IKE proposal.
  - Manage—Opens the Configure IKE Proposals dialog box.
  - IPsec Proposal—Specifies one or more encryption algorithms to use for the IPsec proposal.

Modes

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

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

Adding or Editing a Site-to-Site Tunnel Group

The Add or Edit IPsec Site-to-Site Tunnel Group dialog box lets you specify attributes for the IPsec site-to-site connection that you are adding. In addition, you can select IKE peer and user authentication parameters, configure IKE keepalive monitoring, and select the default group policy.

Fields

- Name—Specifies the name assigned to this tunnel group. For the Edit function, this field is display-only.
- IKE Authentication—Specifies the pre-shared key and Identity certificate parameters to use when authenticating an IKE peer.
  - Pre-shared Key—Specify the value of the pre-shared key for the tunnel group. The maximum length of the pre-shared key is 128 characters.
  - Identity Certificate—Specifies the name of the ID certificate to use for authentication, if available.

  - Manage—Opens the Manage Identity Certificates dialog box, on which you can see the certificates that are already configured, add new certificates, show details for a certificate, and edit or delete a certificate.
  - IKE Peer ID Validation—Specifies whether to check IKE peer ID validation. The default is Required.
- IKE Keepalive ——Enables and configures IKE keepalive monitoring. You can select only one of the following attributes.
  - Disable Keep Alives—Enables or disables IKE keep alives.
  - Monitor Keep Alives—Enables or disables IKE keep alive monitoring. Selecting this option makes available the Confidence Interval and Retry Interval fields.
  - Confidence Interval—Specifies the IKE 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 10 seconds.
  - Retry Interval—Specifies number of seconds to wait between IKE 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.
- Default Group Policy—Select the group policy and client protocols that you want to use as the default for this connection. A VPN group policy is a collection of user-oriented attribute-value pairs that can be stored internally on the device or externally on a RADIUS server. IPsec connections and user accounts refer to the group-policy information.
  - Group Policy—Lists the currently configured group policies. The default value is DfltGrpPolicy.
  - Manage—Opens the Configure Group Policies dialog box, on which you can view the configured group policies and add, edit, or delete group policies from the list.
  - IPsec Protocol—Enables or disables the IPsec protocol for use by this 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
•	—	•	—	—

Crypto Map Entry

In this dialog box, specify crypto parameters for the Connection Profile.

Fields

- **Priority**—A unique priority (1 through 65,543, with 1 the highest priority). 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.

- **Perfect Forward Secrecy**—Ensures that the key for a given IPsec SA was not derived from any other secret (like some other keys). If someone were to break a key, PFS ensures that the attacker would not be able to derive any other key. If you enable PFS, the Diffie-Hellman Group list becomes active.
  - **Diffie-Hellman Group**—An identifier which the two IPsec peers use to derive a shared secret without transmitting it to each other. The choices are Group 1 (768-bits), Group 2 (1024-bits), and Group 5 (1536-bits).
- **Enable NAT-T**— Enables NAT Traversal (NAT-T) for this policy, which lets IPsec peers establish both remote access and LAN-to-LAN connections through a NAT device.
- **Enable Reverse Route Injection**—Provides the ability for static routes to be automatically inserted into the routing process for those networks and hosts that are protected by a remote tunnel endpoint.
- **Security Association Lifetime**—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.

Crypto Map Entry for Static Peer Address

In this dialog box, specify crypto parameters for the Connection Profile when the Peer IP Address is a static address.

Fields

- **Priority**—A unique priority (1 through 65,543, with 1 the highest priority). 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.
- **Perfect Forward Secrecy**—Ensures that the key for a given IPsec SA was not derived from any other secret (like some other keys). If someone were to break a key, PFS ensures that the attacker would not be able to derive any other key. If you enable PFS, the Diffie-Hellman Group list becomes active.
  - **Diffie-Hellman Group**—An identifier which the two IPsec peers use to derive a shared secret without transmitting it to each other. The choices are Group 1 (768-bits), Group 2 (1024-bits), and Group 5 (1536-bits).
- **Enable NAT-T**— Enables NAT Traversal (NAT-T) for this policy, which lets IPsec peers establish both remote access and LAN-to-LAN connections through a NAT device.
- **Enable Reverse Route Injection**—Provides the ability for static routes to be automatically inserted into the routing process for those networks and hosts that are protected by a remote tunnel endpoint.
- **Security Association Lifetime**—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.

- **Static Crypto Map Entry Parameters**—Configure these additional parameters when the Peer IP Address is specified as Static:
  - **Connection Type**—Specify the allowed negotiation as bidirectional, answer-only, or originate-only.
  - **Send ID Cert. Chain**—Enables transmission of the entire certificate chain.
  - **IKE Negotiation Mode**—Sets the mode for exchanging key information for setting up the SAs, Main or Aggressive. It also 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**—An identifier which the two IPsec peers use to derive a shared secret without transmitting it to each other. The choices are Group 1 (768-bits), Group 2 (1024-bits), and Group 5 (1536-bits).

Managing CA Certificates

Clicking Manage under IKE Peer Authentication opens the Manage CA Certificates dialog box. Use this dialog box to view, add, edit, and delete entries on the list of CA certificates available for IKE peer authentication.

The Manage CA Certificates dialog box lists information about currently configured certificates, including information about whom the certificate was issued to, who issued the certificate, when the certificate expires, and usage data.

Fields

- Add or Edit—Opens the Install Certificate dialog box or the Edit Certificate dialog box, which let you specify information about and install a certificate.
- Show Details—Displays detailed information about a certificate that you select in the table.
- Delete—Removes the selected certificate 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
•	—	•	—	—

Install Certificate

Use this dialog box to install a new CA certificate. You can get the certificate in one of the following ways:

- Install from a file by browsing to the certificate file.

- Paste the previously acquired certificate text in PEM format into the box on this dialog box.
- Use SCEP—Specifies the use of the Simple Certificate Enrollment Protocol (SCEP) Add-on for Certificate Services runs on the Windows Server 2003 family. It provides support for the SCEP protocol, which allows Cisco routers and other intermediate network devices to obtain certificates.
  - SCEP URL: http://—Specifies the URL from which to download SCEP information.
  - Retry Period—Specifies the number of minutes that must elapse between SCEP queries.
  - Retry Count—Specifies the maximum number of retries allowed.
- More Options—Opens the Configure Options for CA Certificate 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
•	—	•	—	—

Configure Options for CA Certificate

Use this dialog box to specify details about retrieving CA Certificates for this IPsec remote access connection. The dialog boxes on this dialog box are: Revocation Check, CRL Retrieval Policy, CRL Retrieval Method, OCSP Rules, and Advanced.

Revocation Check Dialog Box

Use this dialog box to specify information about CA Certificate revocation checking.

Fields

- The radio buttons specify whether to check certificates for revocation. The values of these buttons are as follows:
  - Do not check certificates for revocation
  - Check Certificates for revocation
- Revocation Methods area—Lets you specify the method–CRL or OCSP–to use for revocation checking, a nd the order in which to use these methods. You can choose either or both methods.

Add/Edit Remote Access Connections > Advanced > General

Use this dialog box to specify whether to strip the realm and group from the username before passing them to the AAA server, and to specify password management parameters.

Fields

- Strip the realm from username before passing it on to the AAA server—Enables or disables stripping the realm (administrative domain) 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.

The security appliance does not include support for the user@grouppolicy, as the VPN 3000 Concentrator did. Only the L2TP/IPsec client supports the tunnel switching via user@tunnelgroup.

- 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. You can select either to notify the user at login a specific number of days before the password expires or to notify the user only on the day that the password expires. The default is to notify the user 14 days prior to password expiration and every day thereafter until the user changes the password. The range is 1 through 180 days.

Note This does not change the number of days before the password expires, but rather, it enables the notification. If you select this option, you must also specify the number of days.

In either case, and, if the password expires without being changed, the security appliance 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.

This feature requires the use of MS-CHAPv2.

Modes

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

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

Configuring Client Addressing

To specify the client IP address assignment policy and assign address pools to all IPsec and SSL VPN connections, choose Config > Remote Access VPN > Network (Client) Access > IPsec or SSL VPN Connections > Add or Edit > Advanced > Client Addressing. The Add IPsec Remote Access Connection or Add SSL VPN Access Connection opens. Use this dialog box to add address pools and assign them to interfaces, and view, edit, or delete them. The table at the bottom of the dialog box lists the configured interface-specific address pools.

To understand the fields in this dialog box or its descendent dialog boxes, see the sections that follow this one. You can view or change the configuration of address pools and their assignment to interfaces, as follows:

- To view or change the configuration of address pools, click **Add** or **Edit** in the Add IPsec Remote Access Connection or Add SSL VPN Access Connection dialog box. The Assign Address Pools to Interface dialog box opens. This dialog box lets you assign IP address pools to the interfaces configured on the security appliance. Click **Select**. The Select Address Pools dialog box opens. Use this dialog box to view the configuration of address pools. You can change their address pool configuration as follows:
  - To add an address pool to the security appliance, choose **Add**. The Add IP Pool dialog box opens.
  - To change the configuration of an address pool on the security appliance, choose **Edit**. The Edit IP Pool dialog box opens if the addresses in the pool are not in use.

    Note You cannot modify an address pool if it is already in use. If you click **Edit** and the address pool is in use, ASDM displays an error message and lists the connection names and usernames that are using the addresses in the pool.

  - To remove address pool on the security appliance, select the entry in the table and click **Delete**.

Note You cannot remove an address pool if it is already in use. If you click **Delete** and the address pool is in use, ASDM displays an error message and lists the connection names that are using the addresses in the pool.

- To assign address pools to an interface, click **Add** in the Add IPsec Remote Access Connection or Add SSL VPN Access Connection dialog box. The Assign Address Pools to Interface dialog box opens. Select the interface to be assigned an address pool. Click **Select** next to the Address Pools field. The Select Address Pools dialog box opens. Double-click each unassigned pool you want to assign to the interface or choose each unassigned pool and click **Assign**. The adjacent field displays the list of pool assignments. Click OK to populate the Address Pools field with the names of these address pools, then **OK** again to complete the configuration of the assignment.
- To change the address pools assigned to an interface, double-click the interface, or choose the interface in the Add IPsec Remote Access Connection or Add SSL VPN Access Connection dialog box and click Edit. The Assign Address Pools to Interface dialog box opens. To remove address pools, double-click each pool name and press the Delete button on the keyboard. Click **Select** next to the Address Pools field if you want to assign additional fields to the interface. The Select Address Pools dialog box opens. Note that the Assign field displays the address pool names that remained assigned to the interface. Double-click each unassigned pool you want to add to the interface. The Assign field updates the list of pool assignments. Click **OK** to revise the Address Pools field with the names of these address pools, then **OK** again to complete the configuration of the assignment.
- To remove an entry from the Add IPsec Remote Access Connection or Add SSL VPN Access Connection dialog box, choose the entry and click **Delete**.

The Add IPsec Remote Access Connection and Add SSL VPN Access Connection dialog boxes and their descendent dialog boxes are identical. Use the following sections to understand or assign values to the fields in these dialog boxes:

- Add IPsec Remote Access Connection and Add SSL VPN Access Connection
- Assign Address Pools to Interface
- Select Address Pools
- Add or Edit IP Pool
- Add or Edit IP Pool

Add IPsec Remote Access Connection and Add SSL VPN Access Connection

To access the Add IPsec Remote Access Connection and Add SSL VPN Access Connection dialog boxes, choose Config > Remote Access VPN > Network (Client) Access > IPsec or SSL VPN Connections > Add or Edit > Advanced > Client Addressing.

Fields

Use the following descriptions to assign values to the fields in this dialog box:

- Global Client Address Assignment Policy—Configures a policy that affects all IPsec and SSL VPN Client connections (including AnyConnect client connections). The security appliance uses the selected sources in order, until it finds an address:
  - Use authentication server—Specifies that the security appliance should attempt to use the authentication server as the source for a client address.
  - Use DHCP—Specifies that the security appliance should attempt to use DHCP as the source for a client address.

  - Use address pool—Specifies that the security appliance should attempt to use address pools as the source for a client address.
- Interface-Specific Address Pools—Lists the configured interface-specific address pools.

Modes

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

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

Assign Address Pools to Interface

Use the Assign Address Pools to Interface dialog box to select an interface and assign one or more address pools to that interface. To access this dialog box, choose Config > Remote Access VPN > Network (Client) Access > IPsec or SSL VPN Connections > Add or Edit > Advanced > Client Addressing > Add or Edit.

Fields

Use the following descriptions to assign values to the fields in this dialog box:

- Interface—Select the interface to which you want to assign an address pool. The default is DMZ.
- Address Pools—Specify an address pool to assign to the specified interface.
- Select—Opens the Select Address Pools dialog box, in which you can select one or more address pools to assign to this interface. Your selection appears in the Address Pools field of the Assign Address Pools to Interface 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
•	—	•	—	—

Select Address Pools

The Select Address Pools dialog box shows the pool name, starting and ending addresses, and subnet mask of address pools available for client address assignment and lets you add, edit, or delete entries from that list. To access this dialog box, choose Config > Remote Access VPN > Network (Client) Access > IPsec or SSL VPN Connections > Add or Edit > Advanced > Client Addressing > Add or Edit > Select.

Fields

Use the following descriptions to assign values to the fields in this dialog box:

- Add—Opens the Add IP Pool dialog box, on which you can configure a new IP address pool.
- Edit—Opens the Edit IP Pool dialog box, on which you can modify a selected IP address pool.
- Delete—Removes the selected address pool. There is no confirmation or undo.
- Assign—Displays the address pool names that remained assigned to the interface. Double-click each unassigned pool you want to add to the interface. The Assign field updates the list of pool assignments.

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 IP Pool

The Add or Edit IP Pool dialog box lets you specify or modify a range of IP addresses for client address assignment. To access this dialog box, choose Config > Remote Access VPN > Network (Client) Access > IPsec or SSL VPN Connections > Add or Edit > Advanced > Client Addressing > Add or Edit > Select > Add or Edit.

Fields

Use the following descriptions to assign values to the fields in this dialog box:

- Name—Specifies the name assigned to the IP address pool.
- Starting IP Address—Specifies the first IP address in the pool.
- Ending IP Address—Specifies the last IP address in the pool.
- Subnet Mask—Selects the subnet mask to apply to the addresses in the 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 Connection Profile > General > Authentication

This dialog box is available for IPsec on Remote Access and Site-to-Site tunnel groups. The settings on this dialog box apply to the tunnel group globally across the security appliance. To set authentication server group settings per interface, click Advanced. This dialog box 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, click Advanced.
- 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.

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 SSL VPN Connection > General > Authorization

The settings on this dialog box apply to the connection (tunnel group) globally across the security appliance. This dialog box lets you configure the following attributes:

- Authorization Server Group—Lists the available authorization server groups, including the LOCAL group. 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 SSL VPN Connections > Advanced > Accounting

The settings on this dialog box apply to the connection (tunnel group) globally across the security appliance. This dialog box 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.
- Manage—Opens the Configure AAA Server Groups 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/Edit Tunnel Group > General > Client Address Assignment

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 dialog box > General > Client Address Assignment dialog box, 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, click Advanced.

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 > Advanced

The Add or Edit Tunnel Group dialog box, General, Advanced dialog box, 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

On the Add or Edit Tunnel Group dialog box for IPsec for Remote Access, the IPsec dialog box 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 dialog box)

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 dialog boxes-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 for Site-to-Site VPN

The Add or Edit Tunnel Group dialog box lets you configure or edit tunnel group parameters for this Site-to-Site connection profile.

Fields

- Certificate Settings—Sets the following certificate chain and IKE peer validation attributes:
  - Send certificate chain—Enables or disables sending the entire certificate chain. This action includes the root certificate and any subordinate CA certificates in the transmission.
  - IKE Peer ID Validation—Selects whether IKE peer ID validation is ignored, required, or checked only if supported by a certificate.
- IKE Keep Alive—Enables and configures IKE (ISAKMP) keepalive monitoring.
  - Disable Keepalives—Enables or disables IKE keep alives.
  - Monitor Keepalives—Enables or disables IKE keep alive monitoring. Selecting this option makes available the Confidence Interval and Retry Interval fields.
  - Confidence Interval—Specifies the IKE keepalive 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 IKE keepalive retries. The default is 2 seconds.
  - Head end will never initiate keepalive monitoring—Specifies that the central-site security appliance never initiates keepalive monitoring.
- Default Group Policy—Specifies the following group-policy attributes:
  - Group Policy—Selects a group policy to use as the default group policy. The default value is DfltGrpPolicy.

- Manage—Opens the Configure Group Policies dialog box.
- IPsec Protocol—Enables or disables the use of the IPsec protocol for this connection profile.

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

On the Add or Edit Tunnel Group dialog box for a IPsec remote access tunnel group, the PPP dialog box lets you configure or edit the authentication protocols permitted of a PPP connection. This dialog box 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 > Basic

On the Add or Edit Tunnel Group dialog box for Site-to-Site Remote Access, the General, Basic dialog box 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 dialog box, the General dialog box 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 dialog box.
- 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.

The security appliance does not include support for the user@grouppolicy, as the VPN 3000 Concentrator did. Only the L2TP/IPsec client supports the tunnel switching via user@tunnelgroup.

- 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

The Add or Edit Tunnel Group dialog box for IPsec for Site-to-Site access, IPsec dialog box, lets you configure or edit IPsec Site-to-Site-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 dialog box.
- 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 dialog box) 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
•	—	•	—	—

Clientless SSL VPN Access > Connection Profiles > Add/Edit > General > Basic

The Add or Edit pane, General, Basic dialog box 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 dialog box, the General dialog box displays the name and type of the selected tunnel group. All other functions are the same as for the Add Tunnel Group dialog box.

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 dialog box.
- Group Policy—Lists the currently configured group policies. The default value is the default group policy, DfltGrpPolicy.
- Strip the realm —Not available for Clientless SSL VPN.
- Strip the group —Not available or Clientless SSL VPN.
- 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
•	—	•	—	—

Configuring Internal Group Policy IPsec Client Attributes

Use this dialog box to specify whether to strip the realm and group from the username before passing them to the AAA server, and to specify password management options.

Fields

- Strip the realm from username before passing it on to the AAA server—Enables or disables stripping the realm (administrative domain) 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.

The security appliance does not include support for the user@grouppolicy, as the VPN 3000 Concentrator did. Only the L2TP/IPsec client supports the tunnel switching via user@tunnelgroup.

- 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. You can select either to notify the user at login a specific number of days before the password expires or to notify the user only on the day that the password expires. The default is to notify the user 14 days prior to password expiration and every day thereafter until the user changes the password. The range is 1 through 180 days.

Note This does not change the number of days before the password expires, but rather, it enables the notification. If you select this option, you must also specify the number of days.

In either case, and, if the password expires without being changed, the security appliance 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.

Modes

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

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

Configuring Client Addressing for SSL VPN Connections

Use this dialog box to specify the global client address assignment policy and to configure interface-specific address pools. You can also add, edit, or delete interface-specific address pools using this dialog box. The table at the bottom of the dialog box lists the configured interface-specific address pools.

Fields

- Global Client Address Assignment Policy—Configures a policy that affects all IPsec and SSL VPN Client connections (including AnyConnect client connections). The security appliance uses the selected sources in order, until it finds an address:
  - Use authentication server—Specifies that the security appliance should attempt to use the authentication server as the source for a client address.
  - Use DHCP—Specifies that the security appliance should attempt to use DHCP as the source for a client address.
  - Use address pool—Specifies that the security appliance should attempt to use address pools as the source for a client address.
- Interface-Specific Address Pools—Lists the configured interface-specific address pools.
- Add—Opens the Assign Address Pools to Interface dialog box, on which you can select an interface and select an address pool to assign.
- Edit—Opens the Assign Address Pools to Interface dialog box with the interface and address pool fields filled in.
- Delete—Deletes the selected interface-specific address pool. 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
•	—	•	—	—

Assign Address Pools to Interface

Use this dialog box to select an interface and assign one or more address pools to that interface.

Fields

- Interface—Select the interface to which you want to assign an address pool. The default is DMZ.
- Address Pools—Specify an address pool to assign to the specified interface.
- Select—Opens the Select Address Pools dialog box, in which you can select one or more address pools to assign to this interface. Your selection appears in the Address Pools field of the Assign Address Pools to Interface dialog box.

Select Address Pools

The Select Address Pools dialog box shows the pool name, starting and ending addresses, and subnet mask of address pools available for client address assignment and lets you add, edit, or delete entries from that list.

Fields

- Add—Opens the Add IP Pool dialog box, on which you can configure a new IP address pool.
- Edit—Opens the Edit IP Pool dialog box, on which you can modify a selected IP address pool.
- Delete—Removes the selected address pool. There is no confirmation or undo.
- Assign—Displays the address pool names that remained assigned to the interface. Double-click each unassigned pool you want to add to the interface. The Assign field updates the list of pool assignments.

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 an IP Address Pool

Configures or modifies an IP address pool.

Fields

- Name—Specifies the name assigned to the IP address pool.
- Starting IP Address—Specifies the first IP address in the pool.
- Ending IP Address—Specifies the last IP address in the pool.
- Subnet Mask—Selects the subnet mask to apply to the addresses in the pool.

Modes

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

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

Authenticating SSL VPN Connections

The SSL VPN Connections > Advanced > Authentication dialog box lets you configure authentication attributes for SSL VPN connections.

System Options

The System Options pane 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.
- Preserve stateful VPN flows when tunnel drops for Network-Extension Mode (NEM)—Enables or disables preserving IPsec tunneled flows in Network-Extension Mode. With the persistent IPsec tunneled flows feature enabled, as long as the tunnel is recreated within the timeout dialog box, data continues flowing successfully because the security appliance still has access to the state information. This option is disabled by default.

Note Tunneled TCP flows are not dropped, so they rely on the TCP timeout for cleanup. However, if the timeout is disabled for a particular tunneled flow, that flow remains in the system until being cleared manually or by other means (for example, by a TCP RST from the peer).

Modes

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

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

Configuring SSL VPN Connections, Advanced

The advanced options include configuring split tunneling, IE browser proxy, and group-policy related attributes for SSL VPN/AnyConnect clients and IPsec clients.

Configuring Split Tunneling

Split tunneling lets you specify that certain data traffic is encrypted ("goes through the tunnel"), while the remainder is sent in the clear (unencrypted). Split-tunneling network lists distinguish networks that require traffic to go through the tunnel from those that do not require tunneling. the security appliance makes split-tunneling decisions based on a network list, which is an ACL consisting of a list of addresses on the private network.

Fields

- DNS Names—Specify one or more DNS names to which this policy applies.
- Policy—Selects the split-tunneling policy, specifying whether to include or exclude from the tunnel the indicated network lists. If you do not select Inherit, the default is Exclude Network List Below.
- Network List—Selects the networks to which to apply the split-tunneling policy. If you do not select Inherit, the default is --None--.
- Manage—Opens the ACL Manager dialog box, in which you can configure access control lists to use as network lists.
- Intercept DHCP Configuration Message from Microsoft Clients—Reveals additional parameters specific to DHCP Intercept. DCHP Intercept lets Microsoft XP clients use split-tunneling with the security appliance. For Windows clients prior to XP, DHCP Intercept provides the domain name and subnet mask.
  - Intercept—Specifies whether to allow the DHCP Intercept to occur. If you do not select, Inherit, the default setting is No.
  - Subnet Mask—Selects the subnet mask to use.

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.
- SSL Certificate Port: Specify the security appliance port to be used for SSL Authorization. The default is port 80.
- 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

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 dialog box 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 dialog box 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

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
•	—	•	—	—

AnyConnect Essentials

AnyConnect Essentials is a separately licensed SSL VPN client, entirely configured on the security appliance, that provides the full AnyConnect capability, with the following exceptions:

- No CSD (including HostScan/Vault/Cache Cleaner)
- No clientless SSL VPN
- Optional Windows Mobile Support (requires AnyConnect for Windows Mobile license)

The AnyConnect Essentials client provides remote end users running Microsoft Windows Vista, Windows Mobile, Windows XP or Windows 2000, Linux, or Macintosh OS X, with the benefits of a Cisco SSL VPN client.

To enable AnyConnect Essentials, check the **Enable AnyConnect Essentials** check box on the AnyConnect Essentials pane, which appears only if the AnyConnect Essentials license is installed on the security appliance.

When AnyConnect Essentials is enabled, AnyConnect clients use Essentials mode, and clientless SSL VPN access is disabled. When AnyConnect Essentials is disabled, AnyConnect clients use the full AnyConnect SSL VPN Client.

Note The status information about the AnyConnect Essentials license on the Configuration > Device Management > Licensing > Activation Key pane simply reflects whether the AnyConnect Essentials license is installed. This status is not affected by the setting of the Enable AnyConnect Essentials License check box.

AnyConnect Essentials mode cannot be enabled when active clientless sessions exist to the device. To view SSL VPN session details click the **Monitoring > VPN > VPN Sessions** link in the SSL VPN Sessions section. This opens the Monitoring > VPN > VPN > VPN Statistics > Sessions pane. To see session details, choose **Filter By: Clientless SSL VPN** and click **Filter**. This displays session details.

To see how many clientless SSL VPN sessions are currently active, without showing session details, click **Check Number of Clientless SSL Sessions**. If the SSL VPN session count is zero, you can enable AnyConnect Essentials.

Note Secure Desktop does not work when AnyConnect Essentials is enabled. You can, however, disable AnyConnect Essentials when you enable Secure Desktop.

Modes

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

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

DTLS Settings

Enabling Datagram Transport Layer Security (DTLS) allows the AnyConnect VPN client establishing an SSL VPN connection to use two simultaneous tunnels—an SSL tunnel and a DTLS tunnel. Using DTLS avoids latency and bandwidth problems associated with some SSL connections and improves the performance of real-time applications that are sensitive to packet delays.

If you do not enable DTLS, AnyConnect client users establishing SSL VPN connections connect with an SSL VPN tunnel only.

Fields

- Interface—Displays a list of interfaces on the security appliance.
- DTLS Enabled—Click to enable DTLS connections with the AnyConnect client on the interfaces.
- UDP Port (default 443)—(Optional) Specify a separate UDP port for DTLS connections.

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 Settings

The Cisco AnyConnect VPN client provides secure SSL connections to the security appliance for remote users. The client gives remote users the benefits of an SSL VPN client without the need for network administrators to install and configure clients on remote computers.

Without a previously-installed client, remote users enter the IP address in their browser of an interface configured to accept SSL VPN connections. Unless the security appliance is configured to redirect http:// requests to https://, users must enter the URL in the form https://<*address*>.

After entering the URL, the browser connects to that interface and displays the login screen. If the user satisfies the login and authentication, and the security appliance identifies the user as requiring the client, it downloads the client that matches the operating system of the remote computer. After downloading, the client installs and configures itself, establishes a secure SSL connection and either remains or uninstalls itself (depending on the security appliance configuration) when the connection terminates.

In the case of a previously installed client, when the user authenticates, the security appliance examines the revision of the client, and upgrades the client as necessary.

When the client negotiates an SSL VPN connection with the security appliance, it connects using Transport Layer Security (TLS), and optionally, Datagram Transport Layer Security (DTLS). DTLS avoids latency and bandwidth problems associated with some SSL connections and improves the performance of real-time applications that are sensitive to packet delays.

The AnyConnect client can be downloaded from the security appliance, or it can be installed manually on the remote PC by the system administrator. For more information about installing the client manually, see the *Cisco AnyConnect VPN Client Administrator Guide*.

The security appliance downloads the client based on the group policy or local user policy attributes. You can configure the security appliance to automatically download the client, or you can configure it to prompt the remote user about whether to download the client. In the latter case, if the user does not respond, you can configure the security appliance to either download the client after a timeout period or present the login page.

Fields

- SSL VPN Client Images table—Displays the package files specified as SSL VPN client images, and allows you to establish the order that the security appliance downloads the images to the remote PC.
  - Add—Displays the Add SSL VPN Client Image dialog box, where you can specify a file in flash memory as a client image file, or where you can browse flash memory for a file to specify as a client image. You can also upload a file from a local computer to the flash memory.
  - Replace—Displays the Replace SSL VPN Client Image dialog box, where you can specify a file in flash memory as an client image to replace an image hightlighted in the SSL VPN Client Images table. You can also upload a file from a local computer to the flash memory.
  - Delete—Deletes an image from the table. This does not delete the package file from flash.

- Move Up and Move Down—changes the order in which the security appliance downloads the client images to the remote PC. It downloads the image at the top of the table first. Therefore, you should move the image used by the most commonly-encountered operating system to the top.

- SSL VPN Client Profiles table—Displays the XML files specified as SSL VPN client profiles. These profiles display host information in the AnyConnect VPN Client user interface.
  - Add—Displays the Add SSL VPN Client Profiles dialog box, where you can specify a file in flash memory as a profile, or where you can browse flash memory for a file to specify as a profile. You can also upload a file from a local computer to the flash memory.
  - Edit—Displays the Edit SSL VPN Client Profiles dialog box, where you can specify a file in flash memory as a profile to replace a profile highlighted in the SSL VPN Client Profiles table. You can also upload a file from a local computer to the flash memory.
  - Delete—Deletes a profile from the table. This does not delete the XML file from flash.
- Cache File System—The security appliance expands SSL VPN client and CSD images in cache memory. Adjust the size of cache memory to ensure the images have enough space to expand.

Modes

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

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

Add/Replace SSL VPN Client Image

In this pane, you can specify a filename for a file on the security appliance flash memory that you want to add as an SSL VPN client image, or to replace an image already listed in the table. You can also browse the flash memory for a file to identify, or you can upload a file from a local computer.

Fields

- Flash SVC Image—Specify the file in flash memory that you want to identify as an SSL VPN client image.
- Browse Flash—Displays the Browse Flash dialog box where you can view all the files on flash memory.
- Upload—Displays the Upload Image dialog box where you can upload a file from a local PC that you want to identify as an client image.
- Regular expression to match user-agent—Specifies a string that the security appliance uses to match against the User-Agent string passed by the browser. For mobile users, you can decrease the connection time of the mobile device by using the feature. When the browser connects to the security appliance, it includes the User-Agent string in the HTTP header. When the security appliance receives the string, if the string matches an expression configured for an image, it immediately downloads that image without testing the other client images.

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

In this pane, 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 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 image.
- Browse Local Files—Displays the Select File Path dialog box where you can view all the files on local computer and where you can select a file to identify as a client 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 image.
- Browse Flash—Displays the Browse Flash Dialog dialog box where you can view all the files on flash memory of the security appliance and where you can choose a file to identify as a client image.
- Upload File—Initiates the file upload.

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 SSL VPN Client Profiles

In this pane 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 profile. These profiles display host information in the AnyConnect VPN client user interface. You can also browse the local computer or the flash memory of the security appliance for a file to identify.

Fields

- Profile Name—Associates a name with the XML file that appears in the table. Provide any name that makes it easy for you to remember the hosts identified in the XML profile file.

- Profile Package—Identifies the filename of the file in flash memory on the local computer that you want to identify as an SSL VPN client profile.
- Browse Flash—Displays the Browse Flash Dialog dialog box where you can view all the files on flash memory of the security appliance and where you can choose a file to identify as a profile.
- Upload File—Initiates the file upload.

Modes

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

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

Upload Package

In this pane, 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 profile. 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 profile.
- Browse Local Files—Displays the Select File Path dialog box where you can view all the files on local computer and where you can select a file to identify as a client profile.
- 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 client profile.
- Browse Flash—Displays the Browse Flash dialog box where you can view all the files on flash memory of the security appliance and where you can choose a file to identify as a client profile.
- Upload File—Initiates the file upload.

Modes

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

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

Bypass Interface Access List

You can require an access rule to apply to the local IP addresses by unchecking this check box. The access rule applies to the local IP address, and not to the original client IP address used before the VPN packet was decrypted.

- 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 check box 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.)

CHAPTER 37

Configuring Dynamic Access Policies

This chapter describes how to configure dynamic access policies. It includes the following sections.

- Understanding VPN Access Policies
- Add/Edit Dynamic Access Policies
- Add/Edit AAA Attributes
- Retrieving Active Directory Groups
- Add/Edit Endpoint Attributes
- Operator for Endpoint Category
- DAP Examples

Understanding VPN Access Policies

VPN gateways operate in dynamic environments. Multiple variables can affect each VPN connection, for example, intranet configurations that frequently change, the various roles each user may inhabit within an organization, and logins from remote access sites with different configurations and levels of security. The task of authorizing users is much more complicated in a VPN environment than it is in a network with a static configuration.

Dynamic access policies (DAP) on the security appliance let you configure authorization that addresses these many variables. You create a dynamic access policy by setting a collection of access control attributes that you associate with a specific user tunnel or session. These attributes address issues of multiple group membership and endpoint security. That is, the security appliance grants access to a particular user for a particular session based on the policies you define. It generates a DAP at the time the user connects by selecting and/or aggregating attributes from one or more DAP records. It selects these DAP records based on the endpoint security information of the remote device and the AAA authorization information for the authenticated user. It then applies the DAP record to the user tunnel or session.

The DAP system includes the following components that require your attention:

- DAP Selection Configuration File—A text file containing criteria that the security appliance uses for selecting and applying DAP records during session establishment. Stored on the security appliance. You can use ASDM to modify it and upload it to the security appliance in XML data format. DAP selection configuration files include all of the attributes that you configure. These can include AAA attributes, endpoint attributes, and access policies as configured in network and web-type ACL filter, port forwarding and URL lists,

- DfltAccess Policy—Always the last entry in the DAP summary table, always with a priority of 0. You can configure Access Policy attributes for the default access policy, but it does not contain—and you cannot configure—AAA or endpoint attributes. You cannot delete the DfltAccessPolicy, and it must be the last entry in the summary table.

For more information about Dynamic Access Policies, see the following sections.

- DAP Support for Remote Access Connection Types
- DAP and AAA
- DAP and Endpoint Security
- DAP Connection Sequence
- Test Dynamic Access Policies
- DAP Examples

You can also refer to the *Dynamic Access Deployment Guide* (http://supportwiki.cisco.com/ViewWiki/index.php/ASA_8.x_Dynamic_Access_Policies_%28DAP%29_Deployment_Guide) for more information.

Configuring Dynamic Access Policies

To configure dynamic access policies, in the Configuration > Remote Access VPN > Network (Client) Access or Clientless SSL VPN Access > Dynamic Access Policies pane in ASDM, perform the following steps:

Step 1 To include certain antivirus, antispyware, or personal firewall endpoint attributes, click the CSD configuration link near the top of the pane. Then enable Cisco Secure Desktop *and* Host Scan extensions. This link does not display if you have previously enabled both of these features.

If you enable Cisco Secure Desktop, but do not enable Host Scan extensions, when you apply your changes ASDM includes a link to enable Host Scan configuration.

Step 2 To create a new dynamic access policy, click **Add**. To modify an existing policy, click **Edit**.

Step 3 To test already configured polices, click **Test Dynamic Access Policies**.

Fields

- Priority—Displays the priority of the DAP record. The security appliance uses this value to logically sequence the access lists when aggregating the network and web-type ACLs from multiple DAP records. The security appliance orders the records from highest to lowest priority number, with lowest at the bottom of the table. Higher numbers have a higher priority, that is a DAP record with a value of 4 has a higher priority than a record with a value of 2. You cannot manually sort them.
- Name—Displays the name of the DAP record.
- Network ACL List—Displays the name of the firewall access list that applies to the session.
- Web-Type ACL List—Displays the name of the SSL VPN access list that applies to the session.
- Description—Describes the purpose of the DAP record.
- Test Dynamic Access Policies button—Click to test already configured DAP records.

Modes

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

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

DAP Support for Remote Access Connection Types

The DAP system supports the following remote access methods:

- IPsec VPN
- Clientless (browser-based) SSLVPN
- Cisco AnyConnect SSL VPN
- PIX cut-through proxy (posture assessment not available)

DAP and AAA

DAP complements AAA services. It provides a limited set of authorization attributes that can override those AAA provides. The security appliance selects DAP records based on the AAA authorization information for the user and posture assessment information for the session. The security appliance can select multiple DAP records depending on this information, which it then aggregates to create DAP authorization attributes.

You can specify AAA attributes from the Cisco AAA attribute hierarchy, or from the full set of response attributes that the security appliance receives from a RADIUS or LDAP server. For more information about DAP and AAA, see the section, Add/Edit AAA Attributes.

AAA Attribute Definitions

Table 37-1 defines the AAA selection attribute names that are available for DAP use. The Attribute Name field shows you how to enter each attribute name in a Lua logical expression, which you might do in the Advanced section of the Add/Edit Dynamic Access Policy pane.

Table 37-1 AAA Selection Attributes for DAP Use

Attribute Type	Attribute Name	Source	Value	Max String Length	Description
Cisco	aaa.cisco.grouppolicy	AAA	string	64	Group policy name on the security appliance or sent from a Radius/LDAP server as the IETF-CLass (25) attribute
	aaa.cisco.ipaddress	AAA	number	-	Assigned IP address for full tunnel VPN clients (IPsec, L2TP/IPsec, SSL VPN AnyConnect)
	aaa.cisco.tunnelgroup	AAA	string	64	Connection profile (tunnel group) name
	aaa.cisco.username	AAA	string	64	Name of the authenticated user (applies if using Local authentication/authorization)
LDAP	aaa.ldap.<*label*>	LDAP	string	128	LDAP attribute value pair

Table 37-1 AAA Selection Attributes for DAP Use (continued)

RADIUS	aaa.radius.<number>	RADIUS	string	128	Radius attribute value pair
See Security Appliance Supported RADIUS Attributes and Values for a table that lists RADIUS attributes that the security appliance supports.					

DAP and Endpoint Security

The security appliance obtains endpoint security attributes by using posture assessment methods that you configure. These include Cisco Secure Desktop and NAC. For details, see the Cisco Secure Desktop section of ASDM. Table 37-2 identifies each of the remote access protocols DAP supports, the posture assessment tools available for that method, and the information that tool provides.

Table 37-2 DAP Posture Assessment

Remote Access Protocol	Cisco Secure Desktop	Host Scan	NAC	Cisco NAC Appliance
	Returns files information, registry key values, running processes, operating system	Returns antivirus, antispyware, and personal firewall software information	Returns NAC status	Returns VLAN Type and VLAN IDs
IPsec VPN	—[1]	—	X	X
Cisco AnyConnect VPN	X	X	X	X
Clientless VPN	X	X	—	—
PIX Cut-through Proxy	—	—	—	—

[1] — indicates no; X indicates yes.

Endpoint Attribute Definitions

Table 37-3 defines the endpoint selection attribute names that are available for DAP use.The Attribute Name field shows you how to enter each attribute name in a Lua logical expression, which you might do in the Advanced area in the Add/Edit Dynamic Access Policy pane. The *label* variable identifies the application, filename, process, or registry entry.

Table 37-3 Endpoint Attribute Definitions

Attribute Type	Attribute Name	Source	Value	Max String Length	Description
Antispyware (Requires Cisco Secure Desktop)	endpoint.as["*label*"].exists	Host Scan	true	—	Antispyware program exists
	endpoint.as["*label*"].version		string	32	Version
	endpoint.as["*label*"].description		string	128	Antispyware description
	endpoint.as["*label*"].lastupdate		integer	—	Seconds since update of antispyware definitions

Table 37-3 Endpoint Attribute Definitions (continued)

Attribute Type	Attribute Name	Source	Value	Max String Length	Description
Antivirus (Requires Cisco Secure Desktop)	endpoint.av["*label*"].exists	Host Scan	true	—	Antivirus program exists
	endpoint.av["*label*"].version		string	32	Version
	endpoint.av["*label*"].description		string	128	Antivirus description
	endpoint.av["*label*"].lastupdate		integer	—	Seconds since update of antivirus definitions
Application	endpoint.application.clienttype	Application	string	—	Client type: CLIENTLESS ANYCONNECT IPSEC L2TP
File	endpoint.file["*label*"].exists	Secure Desktop	true	—	The files exists
	endpoint.file["*label*"].lastmodified		integer	—	Seconds since file was last modified
	endpoint.file["*label*"].crc.32		integer	—	CRC32 hash of the file
NAC	endpoint.nac.status	NAC	string	—	User defined status string
Operating System	endpoint.os.version	Secure Desktop	string	32	Operating system
	endpoint.os.servicepack		integer	—	Service pack for Windows
Personal firewall (Requires Secure Desktop)	endpoint.fw["*label*"].exists	Host Scan	true	—	The personal firewall exists
	endpoint.fw["*label*"].version		string	32	Version
	endpoint.fw["*label*"].description		string	128	Personal firewall description
Policy	endpoint.policy.location	Secure Desktop	string	64	Location value from Cisco Secure Desktop
Process	endpoint.process["*label*"].exists	Secure Desktop	true	—	The process exists
	endpoint.process["*label*"].path		string	255	Full path of the process
Registry	endpoint.registry["*label*"].type	Secure Desktop	*dword* *string*	—	dword
	endpoint.registry["*label*"].value		string	255	Value of the registry entry
VLAN	endoint.vlan.type	CNA	string	—	VLAN type: ACCESS AUTH ERROR GUEST QUARANTINE ERROR STATIC TIMEOUT

DAP and AntiVirus, AntiSpyware, and Personal Firewall Programs

The security appliance uses a DAP policy when the user attributes matches the configured AAA and endpoint attributes. The Prelogin Assessment and Host Scan modules of Cisco Secure Desktop return information to the security appliance about the configured endpoint attributes, and the DAP subsystem uses that information to select a DAP record that matches the values of those attributes.

Most, but not all, antivirus, antispyware, and personal firewall programs support active scan, which means that the programs are memory-resident, and therefore always running. Host Scan checks to see if an endpoint has a program installed, and if it is memory-resident as follows:

- If the installed program does not support active scan, Host Scan reports the presence of the software. The DAP system selects DAP records that specify the program.
- If the installed program does support active scan, and active scan is enabled for the program, Host Scan reports the presence of the software. Again the security appliance selects DAP records that specify the program.
- If the installed program does support active scan and active scan is disabled for the program, Host Scan ignores the presence of the software. The security appliance does not select DAP records that specify the program. Further, the output of the **debug trace** command, which includes a lot of information about DAP, does not indicate the program presence, even though it is installed.

DAP Connection Sequence

The following sequence outlines a typical remote access connection establishment.

1. A remote client attempts a VPN connection.
2. The security appliance performs posture assessment, using configured NAC and Cisco Secure Desktop Host Scan values.
3. The security appliance authenticates the user via AAA. The AAA server also returns authorization attributes for the user.
4. The security appliance applies AAA authorization attributes to the session, and establishes the VPN tunnel.
5. The security appliance selects DAP records based on the user AAA authorization information and the session posture assessment information.
6. The security appliance aggregates DAP attributes from the selected DAP records, and they become the DAP policy.
7. The security appliance applies the DAP policy to the session.

Test Dynamic Access Policies

This pane lets you test the retrieval of the set of DAP records configured on the device by specifying authorization attribute value pairs. To specify these pairs, use the Add/Edit buttons associated with the AAA Attribute and Endpoint Attribute tables. The dialogs that display when you click these Add/Edit buttons are similar to those in the Add/Edit AAA Attributes and Add/Edit Endpoint Attributes dialog boxes.

When you enter attribute value pairs and click the “Test” button, the DAP subsystem on the device references these values when evaluating the AAA and endpoint selection attributes for each record. The results display in the “Test Results” text area.

Fields

- Selection Criteria—Determine the AAA and endpoint attributes to test for dynamic access policy retrieval.
- AAA Attributes
  - AAA Attribute—Identifies the AAA attribute.
  - Operation Value—Identifies the attribute as =/!= to the given value.
  - Add/Edit—Click to add or edit a AAA attribute.
- Endpoint Attributes—Identifies the endpoint attribute.
  - Endpoint ID—Provides the endpoint attribute ID.
  - Name/Operation/Value—
  - Add/Edit/Delete—Click to add, edit or delete and endpoint attribute.
- Test Result—Displays the result of the test.
- Test—Click to test the retrieval of the policies you have set.
- Close—Click to close 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
•	•	•	—	—

Add/Edit Dynamic Access Policies

shows teh Add Dynamic Access Policies pane.

Add/Edit DTo add or edit a dynamic access policy, perform the following steps:

Step 1 At the top of the Add/Edit Dynamic Access Policy pane, provide a name (required) and a description (optional) of this dynamic access policy.

Step 2 In the Priority field, set a priority for the dynamic access policy. The security appliance applies access policies in the order you set here, highest number having the highest priority. In the case of DAP records with the same priority setting and conflicting ACL rules, the most restrictive rule applies.

Step 3 In the Add/Edit AAA Attributes field, use the ANY/ALL/NONE drop-down list (unlabeled) to choose whether a user must have any, all, or none of the AAA attribute values you configure to use this dynamic access policy.

Step 4 To Set AAA attributes, click **Add/Edit** in the AAA Attributes field.

Step 5 Before you set endpoint attributes, configure CSD Host Scan.

Step 6 To set endpoint security attributes, click **Add/Edit** in the Endpoint ID field.

Step 7 You can create multiple instances of each type of endpoint attribute. For each of these types, you need to decide whether the DAP policy should require that the user have all instances of a type (Match all = AND) or only one of them (Match Any = OR). To set this value for each of the end point attributes, click the **Logical Op.** button.

Step 8 In the Advanced field you can enter one or more logical expressions to set AAA or endpoint attributes other than what is possible in the AAA and Endpoint areas above.

Step 9 To configure network and webtype ACLs, file browsing, file server entry, HTTP proxy, URL entry, port forwarding lists and URL lists, set values in the Access Policy Attributes fields.

Fields

- Policy Name—A string of 4 through 32 characters, no spaces allowed.
- Description—(Optional) Describes the purpose of the DAP record. Maximum 80 characters.
- Priority—Sets the priority of the DAP. The security appliance applies access policies in the order you set here, highest number having the highest priority. Values of 0 to 2147483647 are valid. Default = 0.
- ANY/ALL/NONE drop-down list—Set to require that user authorization attributes match any, all, or none of the values in the AAA attributes you are configuring, as well as satisfying every endpoint attribute. Duplicate entries are not allowed. If you configure a DAP record with no AAA or endpoint attributes, the security appliance always selects it since all selection criteria are satisfied.
- AAA Attributes—Displays the configured AAA attributes.
  - Attribute—Displays the name of the AAA attribute.
  - Operation/Value—=/!=
  - Add/Edit/Delete—Click to add, edit, or delete the highlighted AAA attribute.
- Endpoint Attributes—Displays the configured endpoint attributes
  - Endpoint ID—Identifies endpoint attributes.
  - Name/Operation/Value—Summarizes configured values for each endpoint attribute.
  - Add/Edit/Delete—Click to add, edit, or delete the highlighted endpoint attribute.

Note Cisco Secure Desktop provides the security appliance with all endpoint attributes except Application and NAC. To configure all other endpoint attributes, you must first enable Cisco Secure Desktop, and configure the relevant endpoint attributes there as well.

  - Logical Op.—You can create multiple instances of each type of endpoint attribute. Click to configure whether the DAP policy should require that the user have all instances of a type (Match all = AND) or only one of them (Match Any = OR). Be aware that for some endpoint attributes, for example OS, it can never happen that a user would have more than one instance of the attribute.
  - Advanced—Click to set additional attributes for the dynamic access policy. Be aware that this is an advanced feature that requires knowledge of Lua.
  - AND/OR—Click to define the relationship between the basic selection rules and the logical expressions you enter here, that is, whether the new attributes add to or substitute for the AAA and endpoint attributes already set. The default is AND.

   - Logical Expressions—You can configure multiple instances of each type of endpoint attribute. Enter free-form Lua text that defines new AAA and/or endpoint selection attributes. ASDM does not validate text that you enter here; it just copies this text to the DAP XML file, and the security appliance processes it, discarding any expressions it cannot parse.
   - Guide—Click to display online help for creating these logical operations.
- Access Policy Attributes—These tabs let you set attributes for network and webtype ACL filters, file access, HTTP proxy, URL entry and lists, port forwarding, and clientless SSL VPN access methods. Attribute values that you configure here override authorization values in the AAA system, including those in existing user, group, tunnel group, and default group records.
- Action Tab
   - Action—Specifies special processing to apply to a specific connection or session.
   - Continue—(Default) Click to apply access policy attributes to the session.
   - Terminate—Click to terminate the session.
   - User Message—Enter a text message to display on the portal page when this DAP record is selected. Maximum 128 characters. A user message displays as a yellow orb. When a user logs on it blinks three times to attract attention, and then it is still. If several DAP records are selected, and each of them has a user message, all of the user messages display.

Note You can include in such messages URLs or other embedded text, which require that you use the correct HTML tags.

For example: All contractors please read <a href='http://wwwin.abc.com/procedure.html'> Instructions</a> for the procedure to upgrade your antivirus software.

- Network ACL Filters Tab—Lets you select and configure network ACLs to apply to this DAP record. An ACL for DAP can contain permit or deny rules, but not both. If an ACL contains both permit and deny rules, the security appliance rejects it.
   - Network ACL drop-down list—Select already configured network ACLs to add to this DAP record. Only ACLs having all permit or all deny rules are eligible, and these are the only ACLs that display here.
   - Manage...—Click to add, edit, and delete network ACLs.
   - Network ACL list—Displays the network ACLs for this DAP record.
   - Add—Click to add the selected network ACL from the drop-down list to the Network ACLs list on the right.
   - Delete—Click to delete a highlighted network ACL from the Network ACLs list. You cannot delete an ACL from the security appliance unless you first delete it from DAP records.
- Web-Type ACL Filters Tab—Lets you select and configure web-type ACLs to apply to this DAP record. An ACL for DAP can contain only permit or deny rules. If an ACL contains both permit and deny rules, the security appliance rejects it.
   - Web-Type ACL drop-down list—Select already configured web-type ACLs to add to this DAP record. Only ACLs having all permit or all deny rules are eligible, and these are the only ACLs that display here.
   - Manage...—Click to add, edit, and delete web-type ACLs.
   - Web-Type ACL list—Displays the web-type ACLs for this DAP record.

  - Add—Click to add the selected web-type ACL from the drop-down list to the Web-Type ACLs list on the right.
  - Delete—Click to delete a web-type ACL from the Web-Type ACLs list. You cannot delete an ACL from the security appliance unless you first delete it from DAP records.
- Functions Tab—Lets you configure file server entry and browsing, HTTP proxy, and URL entry for the DAP record.
  - File Server Browsing—Enables or disables CIFS browsing for file servers or shared features.

Note Browsing requires NBNS (Master Browser or WINS). If that fails or is not configured, we use DNS.

The CIFS browse feature does not support internationalization.

  - File Server Entry—Lets or prohibits a user from entering file server paths and names on the portal page. When enabled, places the file server entry drawer on the portal page. Users can enter pathnames to Windows files directly. They 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.
  - HTTP Proxy—Affects 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.
  - URL Entry—Allows or prevents a user from entering HTTP/HTTPS URLs on the portal page. If this feature is enabled, users can enter web addresses in the URL entry box, and use clientless SSL VPN to access those websites.

Using SSL VPN does not ensure that communication with every site is secure. SSL VPN ensures the security of data transmission between the remote user 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 clientless VPN connection, the security appliance acts as a proxy between the end user web browser and target web servers. When a user connects to an SSL-enabled web server, the security appliance establishes a secure connection and validates the server SSL certificate. The end user browser never receives the presented certificate, so therefore cannot examine and validate the certificate. The current implementation of SSL VPN does not permit communication with sites that present expired certificates. Neither does the security appliance perform trusted CA certificate validation. Therefore, users cannot analyze the certificate an SSL-enabled web-server presents before communicating with it.

To limit Internet access for users, choose Disable for the URL Entry field. This prevents SSL VPN users from surfing the web during a clientless VPN connection.

  - Unchanged—(default) Click to use values from the group policy that applies to this session.
  - Enable/Disable—Click to enable or disable the feature.

  - Auto-start—Click to enable HTTP proxy and to have the DAP record automatically start the applets associated with these features.

- Port Forwarding Lists Tab—Lets you select and configure port forwarding lists for user sessions.

  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.

Caution Make sure Sun Microsystems Java Runtime Environment (JRE) 1.4+ is installed on the remote computers to support port forwarding (application access) and digital certificates.

  - Port Forwarding—Select an option for the port forwarding lists that apply to this DAP record. The other attributes in this field are enabled only when you set Port Forwarding to Enable or Auto-start.
  - Unchanged—Click to remove the attributes from the running configuration.
  - Enable/Disable—Click to enable or disable port forwarding.
  - Auto-start—Click to enable port forwarding, and to have the DAP record automatically start the port forwarding applets associated with its port forwarding lists.
  - Port Forwarding List drop-down list—Select already configured port forwarding lists to add to the DAP record.
  - New...—Click to configure new port forwarding lists.
  - Port Forwarding Lists (unlabeled)—Displays the port forwarding lists for the DAP record.
  - Add—Click to add the selected port forwarding list from the drop-down list to the Port Forwarding list on the right.
  - Delete—Click to delete selected port forwarding list from the Port Forwarding list. You cannot delete a port forwarding list from the security appliance unless you first delete it from DAP records.

- URL Lists Tab—Lets you select and configure URL lists for user sessions.
  - Enable URL Lists—Click to enable. When unchecked, no URL lists display in the portal page for the connection.
  - URL List drop-down list—select already configured URL lists to add to the DAP record.
  - Manage...—Click to add, import, export, and delete URL lists.
  - URL Lists (unlabeled)—Displays the URL lists for the DAP record.
  - Add—Click to add the selected URL list from the drop-down list to the URL area on the right.
  - Delete—Click to delete the selected URL list from the URL list area. You cannot delete a URL list from the security appliance unless you first delete it from DAP records.

- Access Method Tab—Lets you configure the type of remote access permitted.
  - Unchanged—Continue with the current remote access method.

- AnyConnect Client—Connect using the Cisco AnyConnect VPN Client.
- Web-Portal—Connect with clientless VPN.
- Both-default-Web-Portal—Connect via either clientless or the AnyConnect client, with a default of clientless.
- Both-default-AnyConnect Client—Connect via either clientless or the AnyConnect client, with a default of AnyConnect.

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 AAA Attributes

To configure AAA attributes as selection criteria for DAP records, in the Add/Edit AAA Attributes dialog box, set the Cisco, LDAP, or RADIUS attributes that you want to use. You can set these attributes either to = or != the value you enter. There is no limit for the number of AAA attributes for each DAP record. For detailed information about AAA attributes, see AAA Attribute Definitions.

Fields

AAA Attributes Type—Use the drop-down list to select Cisco, LDAP or RADIUS attributes:

- Cisco—Refers to user authorization attributes that are stored in the AAA hierarchical model. You can specify a small subset of these attributes for the AAA selection attributes in the DAP record.These include:
  - Group Policy —The group policy name associated with the VPN user session. Can be set locally on the security appliance or sent from a RADIUS/LDAP server as the IETF-Class (25) attribute. Maximum 64 characters.
  - IP Address—The assigned IP address for full tunnel VPN clients (IPsec, L2TP/IPsec, SSL VPN AnyConnect). Does not apply to Clientless SSL VPN, since there is no address assignment for clientless sessions.
  - Connection Profile—The connection or tunnel group name. Maximum 64 characters.
  - Username—The username of the authenticated user. Maximum 64 characters. Applies if you are using Local, RADIUS, LDAP authentication/authorization or any other authentication type (for example, RSA/SDI), NT Domain, etc).
  - =/!=—Equal to/Not equal to.
- LDAP—The LDAP client (security appliance) stores all native LDAP response attribute value pairs in a database associated with the AAA session for the user. The LDAP client writes the response attributes to the database in the order in which it receives them. It discards all subsequent attributes with that name. This scenario might occur when a user record and a group record are both read from the LDAP server. The user record attributes are read first, and always have priority over group record attributes.

To support Active Directory group membership, the AAA LDAP client provides special handling of the LDAP memberOf response attribute. The AD memberOf attribute specifies the DN string of a group record in AD. The name of the group is the first CN value in the DN string. The LDAP client extracts the group name from the DN string and stores it as the AAA memberOf attribute, and in the response attribute database as the LDAP memberOf attribute. If there are additional memberOf attributes in the LDAP response message, then the group name is extracted from those attributes and is combined with the earlier AAA memberOf attribute to form a comma separated string of group names, also updated in the response attribute database.

In the case where the VPN remote access session to an LDAP authentication/authorization server returns the following three Active directory groups (memberOf enumerations):

cn=Engineering,ou=People,dc=company,dc=com

cn=Employees,ou=People,dc=company,dc=com

cn=EastCoastast,ou=People,dc=company,dc=com

the ASA processes three Active Directory groups: Engineering, Employees, and EastCoast which could be used in any combination as aaa.ldap selection criteria.

LDAP attributes consist of an attribute name and attribute value pair in the DAP record. The LDAP attribute name is syntax/case sensitive. If for example you specify LDAP attribute Department instead of what the AD server returns as department, the DAP record will not match based on this attribute setting.

Note To enter multiple values in the Value field, use the semicolon (;) as the delimiter. For example:

eng;sale; cn=Audgen VPN,ou=USERS,o=OAG

- RADIUS—The RADIUS client stores all native RADIUS response attribute value pairs in a database associated with the AAA session for the user. The RADIUS client writes the response attributes to the database in the order in which it receives them. It discards all subsequent attributes with that name. This scenario might occur when a user record and a group record are both read from the RADIUS server. The user record attributes are read first, and always have priority over group record attributes.

  RADIUS attributes consist of an attribute number and attribute value pair in the DAP record. See Security Appliance Supported RADIUS Attributes and Values for a table that lists RADIUS attributes that the security appliance supports.

Note For RADIUS attributes, DAP defines the Attribute ID = 4096 + RADIUS ID.

For example:

The RADIUS attribute "Access Hours" has a Radius ID = 1, therefore DAP attribute value = 4096 + 1 = 4097.

The RADIUS attribute "Member Of" has a Radius ID = 146, therefore DAP attribute value = 4096 + 146 = 4242.

- LDAP and RADIUS attributes include:
  - Attribute ID—Names/numbers the attribute. Maximum 64 characters.

  - Value—The attribute name (LDAP) or number (RADIUS).

    To enter multiple values in the Value field, use the semicolon (;) as the delimiter. For example:

    eng;sale; cn=Audgen VPN,ou=USERS,o=OAG

  - =/!=—Equal to/Not equal to.

- LDAP includes the Get AD Groups button. This button queries the Active Directory LDAP server for the list of groups the user belong to (memberOf enumerations). It retrieves the AD groups using the CLI show-ad-groups command in the background

The **show ad-groups** command applies only to Active Directory servers using LDAP. Use this command to display AD groups that you can use for dynamic access policy AAA selection criteria.

The default time that the security appliance waits for a response from the server is 10 seconds. You can adjust this time using the **group-search-timeout** command in aaa-server host configuration mode.

Note If the Active Directory server has a large number of groups, the output of the **show ad-groups** command might be truncated based on limitations to the amount of data the server can fit into a response packet. To avoid this problem, use the filter option to reduce the number of groups reported 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
•	•	•	—	—

Retrieving Active Directory Groups

You can query an Active Directory server for available AD groups in this pane. This feature applies only to Active Directory servers using LDAP. Use the group information to specify dynamic access policy AAA selection criteria.

You can change the level in the Active Directory hierarchy where the search begins by changing the Group Base DN in the Edit AAA Server pane. You can also change the time that the security appliance waits for a response from the server in the window. To configure these features, choose Configuration > Remote Access VPN > AAA/Local Users > AAA Server Groups > Edit AAA Server.

Note If the Active Directory server has a large number of groups, the list of AD groups retrieved may be truncated based on limitations of the amount of data the server can fit into a response packet. To avoid this problem, use the filter feature to reduce the number of groups reported by the server.

Fields

AD Server Group—The name of the AAA server group to retrieve AD groups.

Filter By—Specify a group or the partial name of a group to reduce the groups displayed.

Group Name—A list of AD groups retrieved from 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 Endpoint Attributes

Endpoint attributes contain information about the endpoint system environment, posture assessment results, and applications. The security appliance dynamically generates a collection of endpoint attributes during session establishment, and stores these attributes in a database associated with the session. There is no limit for the number of endpoint attributes for each DAP record.

Each DAP record specifies the endpoint selection attributes that must be satisfied for the security appliance to select it. The security appliance selects only DAP records that satisfy every condition configured.

For detailed information about Endpoint attributes, see Endpoint Attribute Definitions.

To configure endpoint attributes as selection criteria for DAP records, in the Add/Edit Endpoint Attribute dialog box, set components. These components change according to the attribute type you select.

Fields

- Endpoint Attribute Type—Choose from the drop-down list the endpoint attribute you want to set. Options include Antispyware, Antivirus, Application, File, NAC, Operating System, Personal Firewall, Process, Registry, VLAN, and Priority.

Endpoint attributes include these components, but not all attributes include all components. The following descriptions show (in parentheses) the attributes to which each component applies.

- Exists/Does not exist buttons (Antispyware, Antivirus, Application, File, NAC, Operating System, Personal Firewall, Process, Registry, VLAN, Priority)— Click the appropriate button to indicate whether the selected endpoint attribute and its accompanying qualifiers (fields below the Exists/Does not exist buttons) should be present or not.
- Vendor ID (Antispyware, Antivirus, Personal Firewall)—Identify the application vendor.
- Vendor Description (Antispyware, Antivirus, Personal Firewall)—Provide text that describes the application vendor.
- Version (Antispyware, Antivirus, Personal Firewall)—Identify the version of the application, and specify whether you want the endpoint attribute to be equal to/not equal to that version.
- Last Update (Antispyware, Antivirus, File)—Specify the number of days since the last update. You might want to indicate that an update should occur in less than (<) or more than (>) the number of days you enter here.
- Client Type (Application)—Indicate the type of remote access connection, AnyConnect, Clientless, Cut-through Proxy, IPsec, or L2TP.
- Checksum (File)—Select the file and click the Compute Checksum button to arrive at this value.
- Compute CRC32 Checksum (File)—Use this calculator to determine the checksum value of a file.
- Posture Status (NAC)—Contains the posture token string received from ACS.

- OS Version (Operating System)—Windows (various), MAC, Linux, Pocket PC.
- Service Pack (Operating System)—Identify the service pack for the operating system.
- Endpoint ID (File, Process, Registry)—A string that identifies an endpoint for files, processes or registry entries. DAP uses this ID to match Cisco Secure Desktop host scan attributes for DAP selection. You must configure Host Scan before you configure this attribute. When you configure Host Scan, the configuration displays in this pane, so you can select it, reducing the possibility of errors in typing or syntax.
- Path (Process, Policy)—Configure Host Scan before you configure this attribute. When you configure Host Scan, the configuration displays in this pane, so you can select it, reducing the possibility of errors in typing or syntax.
- Value (Registry)—dword or string
- Caseless (Registry)—Select to disregard case in registry entries.
- VLAN ID (VLAN)—A valid 802.1q number ranging from 1 to 4094
- VLAN Type (VLAN)—Possible values include the following:

ACCESS	Posture assessment passed
STATIC	No posture assessment applied
TIMEOUT	Posture assessment failed due to no response
AUTH	Posture assessment still active
GUEST	Posture assessment passed, switch to guest VLAN
QUARANTINE	Posture assessment failed, switch to quarantine VLAN
ERROR	Posture assessment failed due to fatal error

- Policy (Location)—Enter the Cisco Secure Desktop Microsoft Windows location profile, case sensitive.

Modes

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

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

Guide

This section provides information about constructing logical expressions for AAA or Endpoint attributes. Be aware that doing so requires sophisticated knowledge of Lua (www.lua.org).

In the Advanced field you enter free-form Lua text that represents AAA and/or endpoint selection logical operations. ASDM does not validate text that you enter here; it just copies this text to the DAP policy file, and the security appliance processes it, discarding any expressions it cannot parse.

This option is useful for adding selection criteria other than what is possible in the AAA and endpoint attribute areas above. For example, while you can configure the security appliance to use AAA attributes that satisfy any, all, or none of the specified criteria, endpoint attributes are cumulative, and must all be satisfied. To let the security appliance employ one endpoint attribute or another, you need to create appropriate logical expressions in Lua and enter them here.

- For a list of AAA Selection attributes, including proper name syntax for creating logical expressions, see Table 37-1.
- For a list of endpoint selection attributes, including proper name syntax for creating logical expressions, see Table 37-3.

The following sections provide detailed explanations of creating Lua EVAL expressions, as well as examples.

- Syntax for Creating Lua EVAL Expressions
  - Constructing DAP EVAL Expressions
- The DAP CheckAndMsg Function
  - Checking for a Single Antivirus Program
  - Checking for Antivirus Definitions Within the Last 10 Days
  - Checking for a Hotfix on the User PC
  - Checking for Antivirus Programs
  - Checking for Antivirus Programs and Definitions Older than 1 1/2 Days
- Additional Lua Functions
  - OU-Based Match Example
  - Group Membership Example
  - Antivirus Example
  - Antispyware Example
  - Firewall Example
  - Antivirus, Antispyware, or any Firewall Example
- CheckAndMsg with Custom Function Example
- Further Information on Lua

Syntax for Creating Lua EVAL Expressions

This section provides information about the syntax for creating Lua EVAL expressions.

Note If you must use Advanced mode, we recommend that you use EVAL expressions whenever possible for reasons of clarity, which makes verifying the program straightforward.

EVAL(<*attribute*> , <comparison>, {<*value*> | <*attribute*>}, [<type>])

<attribute>	AAA attribute or an attribute returned from Cisco Secure Desktop, see Table 37-1 and Table 37-3 for attribute definitions
<comparison>	One of the following strings (quotation marks required)

	“EQ”	equal
	“NE”	not equal
	“LT”	less than
	“GT”	greater than
	“LE”	less than or equal
	“GE”	greater than or equal
<value>	A string in quotation marks that contains the value to compare the attribute against	
<type>	One of the following strings (quotation marks required)	
	“string”	case-sensitive string comparison
	“caseless”	case-insensitive string comparison
	“integer”	number comparison, converts string values to numbers
	“hex”	number comparison using hexadecimal values, converts hex string to hex numbers
	“version”	compares versions of the form X.Y.Z. where X, Y, and Z are numbers

Example:

EVAL(endpoint.os.version, “EQ”, “Windows XP”, “string”)

Constructing DAP EVAL Expressions

Study these examples for help in creating logical expressions in Lua.

- This endpoint expression tests for a match on CLIENTLESS OR CVC client types:

```
(EVAL(endpoint.application.clienttype,”EQ”,"CLIENTLESS") or
EVAL(endpoint.application.clienttype, “EQ”,"CVC"))
```

- This endpoint expression tests for Norton Antivirus versions 10.x but excludes 10.5.x:

```
(EVAL(endpoint.av[“NortonAV”].version, “GE”, "10",”version”) and
(EVAL(endpoint.av[“NortonAV”].version,”LT”, "10.5", “version”) or
EVAL(endpoint.av[“NortonAV”].version, “GE”, "10.6", “version”)))
```

The DAP CheckAndMsg Function

CheckAndMsg is a Lua function that you can configure DAP to call. It generates a user message based on a condition.

You use ASDM to configure CheckAndMsg through the Advanced field in DAP. The security appliance displays the message to the user only when the DAP record containing the LUA CheckAndMsg function is selected and results in a clientless SSL VPN or AnyConnect termination.

The syntax of the CheckAndMsg function follows:

```
CheckAndMsg(value, “<message string if value is true>”, “<message string if value if
false>”)
```

Be aware of the following when creating CheckAndMsg functions:

- CheckAndMsg returns the value passed in as its first argument.

- Use the EVAL function as the first argument if you do not want to use string comparison. For example:

  ```
  (CheckAndMsg((EVAL(...)) , "true msg", "false msg"))
  ```

  CheckandMsg returns the result of the EVAL function and the security appliances uses it to determine whether to select the DAP record. If the record is selected and results in termination, the security appliance displays the appropriate message.

Checking for a Single Antivirus Program

This example checks if a single antivirus program, in this case McAfee, is installed on the user PC, and displays a message if it is not.

```
(CheckAndMsg(EVAL(endpoint.av[“McAfeeAV”].exists,"NE","true"),"McAfee AV was not
found on your computer", nil))
```

Checking for Antivirus Definitions Within the Last 10 Days

This example checks antivirus definitions within the last 10 days (864000 sec), in particular the last update of the McAfee AV dat file, and displays a message to a user lacking the appropriate update that they need an antivirus update:

```
((CheckAndMsg(EVAL(endpoint.av[“McAfeeAV”].lastupdate,"GT","864000","integer"),"AV
Update needed! Please wait for the McAfee AV till it loads the latest dat file.",nil) ))
```

Checking for a Hotfix on the User PC

This example checks for a specific hotfix. If a user does not have the hotfix on their PC, a message that it is not installed displays.

```
(not CheckAndMsg(EVAL(endpoint.os.windows.hotfix["KB923414"],"EQ","true"),nil,"The
required hotfix is not installed on your PC."))
```

or you could define it this way (which makes more sense):

```
(CheckAndMsg(EVAL(endpoint.os.windows.hotfix["KB923414"],"NE","true"),"The required hotfix
is not installed on your PC.",nil))
```

You can build the expression in this example because the debug dap trace returns:

```
endpoint.os.windows.hotfix["KB923414"] = "true";
```

Checking for Antivirus Programs

You can configure messages so that the end user is aware of and able to fix problems with missing or not running AVs. As a result, if access is denied, the security appliance collects all messages for the DAP that caused the "terminate" condition and displays them in the browser on the logon page. If access is allowed, the security appliance displays all messages generated in the process of DAP evaluation on the portal page.

The following example shows how to use this feature to check on the Norton Antivirus program.

Step 1 Copy and paste the following Lua expression into the Advanced field of the Add/Edit Dynamic Access Policy pane (click the double arrow on the far right to expand the field).

```
(CheckAndMsg(EVAL(endpoint.av["NortonAV"].exists, "EQ", "false"),"Your Norton AV was found
but the active component of it was not enabled", nil) or
CheckAndMsg(EVAL(endpoint.av["NortonAV"].exists, "NE", "true"),"Norton AV was not found on
your computer", nil) )
```

Step 2 In that same Advanced field, click the **OR** button.

Step 3 In the Access Attributes section below, in the leftmost tab, Action**,** click **Terminate**.

Step 4 Connect from a PC that does not have or has disabled Norton Antivirus.

The expected result is that the connection is not allowed *and* the message appears as a blinking ! point.

Step 5 Click the blinking ! to see the message.

Checking for Antivirus Programs *and* Definitions Older than 1 1/2 Days

This example checks for the presence of the Norton and McAfee antivirus programs, and whether the virus definitions are older than 1 1/2 days (10,000 seconds). If the definitions are older than 1 1/2 days, the security appliance terminates the session with a message and links for remediation. To accomplish this task, perform the following steps.

Step 1 Copy and paste the following Lua expression into the Advanced field of the Add/Edit Dynamic Access Policy pane (click the double arrow on the far right to expand the field):

```
((EVAL(endpoint.av["NortonAV"].exists,"EQ","true","string") and
CheckAndMsg(EVAL(endpoint.av["NortonAV"].lastupdate,"GT","10000",integer"),To
remediate <a href='http://www.symantec.com'>Click this link </a>",nil)) or
(EVAL(endpoint.av["McAfeeAV"].exists,"EQ","true","string") and
CheckAndMsg(EVAL(endpoint.av["McAfeeAV"].lastupdate,"GT","10000",integer"),To
remediate <a href='http://www.mcafee.com'>Click this link</a>",nil))
```

Step 2 In that same Advanced field, click **AND**.

Step 3 In the Access Attributes section below, in leftmost tab, Action**,** click **Terminate**.

Step 4 Connect from a PC that has Norton and McAfee antivirus programs with versions that are older than 1 1/2 days.

The expected result is that the connection is not allowed *and* the message appears as a blinking ! point.

Step 5 Click the blinking ! to see the message and links for remediation.

Additional Lua Functions

When working with dynamic access policies for clientless SSL VPN, you might need additional flexibility of match criteria. For example, you might want to apply a different DAP based on the following:

- Organizational Unit (OU) or other level of the hierarchy for the user object
- Group Name that follows a naming convention but has many possible matches—you might require the ability to use a wildcard on group names.

You can accomplish this flexibility by creating a Lua logical expression in the Advanced section of the DAP pane in ASDM.

OU-Based Match Example

DAP can use many attributes returned from an LDAP server in a logical expression. See the DAP trace section for example output of this, or run a debug dap trace.

The LDAP server returns the user Distinguished Name (DN). This implicitly identifies where in the directory the user object is located. For example, if the user DN is CN=Example User,OU=Admins,dc=cisco,dc=com this user is located in OU=Admins,dc=cisco,dc=com. If all administrators are in this OU (or any container below this level) you can use a logical expression to match on this criteria as follows:

```
assert(function()
   if ( (type(aaa.ldap.distinguishedName) == "string") and
        (string.find(aaa.ldap.distinguishedName, "OU=Admins,dc=cisco,dc=com$") ~= nil) )
then
       return true
   end
   return false
end)()
```

In this example, the string.find function allows for a regular expression. Use the $ at the end of the string to anchor this string to the end of the distinguishedName field.

Group Membership Example

You can create a basic logical expression for pattern matching of AD group membership. Because users can be members of multiple groups, DAP parses the response from the LDAP server into separate entries in a table. You need an advanced function to accomplish the following:

- Compare the memberOf field as a string (in the event the user belongs to only one group).
- Iterate through each returned memberOf field if the returned data is of type "table".

The function we have written and tested for this purpose is shown below. In this example, if a user is a member of any group ending with "-stu" they match this DAP.

```
assert(function()
   local pattern = "-stu$"
   local attribute = aaa.ldap.memberOf
   if ((type(attribute) == "string") and
       (string.find(attribute, pattern) ~= nil)) then
       return true
   elseif (type(attribute) == "table") then
       local k, v
       for k, v in pairs(attribute) do
           if (string.find(v, pattern) ~= nil) then
               return true
           end
       end
   end
   return false
end)()
```

Antivirus Example

The following example uses a custom function to check if CSD detects any antivirus software.

```
assert(function()
    for k,v in pairs(endpoint.av) do
```

```
            if (EVAL(v.exists, "EQ", "true", "string")) then
                return true
            end
        end
        return false
    end)()
```

Antispyware Example

The following example uses a custom function to check if CSD detects any antispyware.

```
assert(function()
    for k,v in pairs(endpoint.as) do
         if (EVAL(v.exists, "EQ", "true", "string")) then
              return true
         end
    end
    return false
end)()
```

Firewall Example

The following example uses a custom function to check if CSD detects a firewall.

```
assert(function()
    for k,v in pairs(endpoint.fw) do
         if (EVAL(v.exists, "EQ", "true", "string")) then
              return true
         end
    end
    return false
end)()
```

Antivirus, Antispyware, *or* any Firewall Example

The following example uses a custom function to check if CSD detects any antivirus, antispyware, or any firewall.

```
assert(function()
       function check(antix)
           if (type(antix) == "table") then
               for k,v in pairs(antix) do
                   if (EVAL(v.exists, "EQ", "true", "string")) then
                       return true
                   end
               end
           end
           return false
       end
       return (check(endpoint.av) or check(endpoint.fw) or check(endpoint.as))
   end)()
```

CheckAndMsg with Custom Function Example

You can use the following function to deny access in the absence of an antivirus program. Use it with a DAP that has Action set to terminate.

```
assert( function()
   for k,v in pairs(endpoint.av) do
     if (EVAL(v.exists, "EQ", "true", "string")) then
         return false
     end
  end
  return CheckAndMsg(true, "Please install antivirus software before connecting.", nil)
end)()
```

If a user lacking an antivirus program attempts to log in, DAP displays the following message:

```
Please install antivirus software before connecting.
```

Further Information on Lua

You can find detailed LUA programming information at http://www.lua.org/manual/5.1/manual.html.

Operator for Endpoint Category

You can configure multiple instances of each type of endpoint. In this pane, set each type of endpoint to require only one instance of a type (Match Any = OR) or to have all instances of a type (Match All = AND).

- If you configure only one instance of an endpoint category, you do not need to set a value.
- For some endpoint attributes, it makes no sense to configure multiple instances. For example, no users have more than one running OS.
- You are configuring the Match Any/Match All operation within each endpoint type.

The security appliance evaluates each type of endpoint attribute, and then performs a logical AND operation on all of the configured endpoints. That is, each user must satisfy the conditions of ALL of the endpoints you configure, as well as the AAA attributes.

DAP Examples

The following sections provide examples of useful dynamic access policies.

- Using DAP to Define Network Resources
- Using DAP to Apply a WebVPN ACL
- Enforcing CSD Checks and Applying Policies via DAP

Using DAP to Define Network Resources

This example shows how to configure dynamic access policies as a method of defining network resources for a user or group. The DAP policy named Trusted_VPN_Access permits clientless and AnyConnect VPN access. The policy named Untrusted_VPN_Access permits only clientless VPN access. Table 37-4 summarizes the configuration of each of these policies.

The ASDM path is Configuration > Remote Access VPN > Clientless SSL VPN Access > Dynamic Access Policies > Add/Edit Dynamic Access Policy > Endpoint

Table 37-4 A Simple DAP Configuration for Network Resources

Attribute	Trusted_VPN_Access	Untrusted_VPN_Access
Endpoint Attribute Type Policy	Trusted	Untrusted
Endpoint Attribute Process	ieexplore.exe	—
Advanced Endpoint Assessment	AntiVirus= McAfee Attribute	
CSD Location	Trusted	Untrusted
LDAP memberOf	Engineering, Managers	Vendors
ACL		Web-Type ACL
Access	**AnyConnect** *and* **Web Portal**	**Web Portal**

Using DAP to Apply a WebVPN ACL

DAP can directly enforce a subset of access policy attributes including Network ACLs (for IPsec and AnyConnect), clientless SSL VPN Web-Type ACLs, URL lists, and Functions. It cannot directly enforce, for example, a banner or the split tunnel list, which the group policy enforces. The Access Policy Attributes tabs in the Add/Edit Dynamic Access Policy pane provide a complete menu of the attributes DAP directly enforces.

Active Directory/LDAP stores user group policy membership as the "memberOf" attribute in the user entry. You can define a DAP such that for a user in AD group (memberOf) = Engineering the security appliance applies a configured Web-Type ACL. To accomplish this task, perform the following steps:

Step 1 Navigate to the Add AAA attributes pane (Configuration > Remote Access VPN > Clientless SSL VPN Access > Dynamic Access Policies > Add/Edit Dynamic Access Policy > AAA Attributes section > Add AAA Attribute).

Step 2 For the AAA Attribute type, use the drop-down menu to choose LDAP.

Step 3 In the Attribute ID field, enter memberOf, exactly as you see it here. Case is important.

Step 4 In the Value field, use the drop-down menu to choose =, and in the adjacent field enter Engineering.

Step 5 In the Access Policy Attributes area of the pane, click the Web-Type ACL Filters tab.

Step 6 Use the Web-Type ACL drop-down menu to select the ACL you want to apply to users in the AD group (memberOf) = Engineering.

Enforcing CSD Checks and Applying Policies via DAP

This example creates a DAP that checks that a user belongs to two specific AD/LDAP groups (Engineering and Employees) and a specific ASA tunnel group. It then applies an ACL to the user.

The ACLs that DAP applies control access to the resources. They override any ACLS defined the group policy on the security appliance. In addition, the security appliance applied the regular AAA group policy inheritance rules and attributes for those that DAP does not define or control, examples being split tunneling lists, banner, and DNS. To accomplish this task, perform the following steps.

Step 1 Navigate to the Add AAA attributes pane (Configuration > Remote Access VPN > Clientless SSL VPN Access > Dynamic Access Policies > Add/Edit Dynamic Access Policy > AAA Attributes section > Add AAA Attribute).

Step 2 For the AAA Attribute type, use the drop-down menu to choose LDAP.

Step 3 In the Attribute ID field, enter memberOf, exactly as you see it here. Case is important.

Step 4 In the Value field, use the drop-down menu to choose =, and in the adjacent field enter Engineering.

Step 5 In the Attribute ID field, enter memberOf, exactly as you see it here. Case is important.

Step 6 In the Value field, use the drop-down menu to select =, and in the adjacent field enter Employees.

Step 7 For the AAA attribute type, use the drop-down menu to choose Cisco.

Step 8 Check the Tunnel group box, use the drop-down menu to choose =, and in the adjacent drop-down list select the appropriate tunnel group (connection policy).

Step 9 In the Network ACL Filters tab of the Access Policy Attributes area, choose the ACLs to apply to users who meet the DAP criteria defined in the previous steps.

CHAPTER 38

Clientless SSL VPN End User Set-up

This section is for the system administrator who sets up Clientless (browser-based) SSL VPN 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 Clientless SSL VPN. This section includes the following topics:

- Requiring Usernames and Passwords
- Communicating Security Tips
- Configuring Remote Systems to Use Clientless SSL VPN Features
- Capturing Clientless SSL VPN Data

Note We assume you have already configured the security appliance for Clientless SSL VPN.

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, Clientless SSL VPN, 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 38-1 lists the type of usernames and passwords that Clientless SSL VPN users might need to know.

Table 38-1 Usernames and Passwords to Give to Clientless SSL VPN 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
Clientless SSL VPN	Access remote network	Starting a Clientless SSL VPN session
File Server	Access remote file server	Using the Clientless SSL VPN file browsing feature to access a remote file server
Corporate Application Login	Access firewall-protected internal server	Using the Clientless SSL VPN web browsing feature to access an internal protected website
Mail Server	Access remote mail server via Clientless SSL VPN	Sending or receiving e-mail messages

Communicating Security Tips

Advise users always to log out from the session. (To log out of Clientless SSL VPN, click the logout icon on the Clientless SSL VPN toolbar or close the browser.)

Advise users that using Clientless SSL VPN does not ensure that communication with every site is secure. Clientless SSL VPN 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 Clientless SSL VPN Features

Table 38-2 includes the following information about setting up remote systems to use Clientless SSL VPN:

- Starting Clientless SSL VPN
- Using the Clientless SSL VPN 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 38-2 also provides information about the following:

- Clientless SSL VPN requirements, by feature
- Clientless SSL VPN 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 features are available to each Clientless SSL VPN user. Table 38-2 organizes information by feature, so you can skip over the information for unavailable features.

Table 38-2 Clientless SSL VPN Remote System Configuration and End User Requirements

Task	Remote System or End User Requirements	Specifications or Use Suggestions
Starting Clientless SSL VPN	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
	Clientless SSL VPN-supported browser	We recommend the following browsers for Clientless SSL VPN. Other browsers might not fully support Clientless SSL VPN features. On Microsoft Windows: • Internet Explorer version 6.0 • Firefox 1.x On Linux: • Firefox 1.x 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 Clientless SSL VPN	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 Clientless SSL VPN is enabled. For example: https://10.89.192.163 or https://cisco.example.com.
	Clientless SSL VPN username and password	
	[Optional] Local printer	Clientless SSL VPN does not support printing from a web browser to a network printer. Printing to a local printer is supported.

Table 38-2 Clientless SSL VPN Remote System Configuration and End User Requirements (continued)

Task	Remote System or End User Requirements	Specifications or Use Suggestions
Using the Floating Toolbar in a Clientless SSL VPN Connection		A floating toolbar is available to simplify the use of Clientless SSL VPN. 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 Clientless SSL VPN session. If you click the **Close** button, the security appliance prompts you to confirm that you want to close the Clientless SSL VPN session. **Tip** TIP: To paste text into a text field, use Ctrl-V. (Right-clicking is disabled on the Clientless SSL VPN toolbar.)
Web Browsing	Usernames and passwords for protected websites	Using Clientless SSL VPN does not ensure that communication with every site is secure. See “Communicating Security Tips.”
		The look and feel of web browsing with Clientless SSL VPN might be different from what users are accustomed to. For example: • The Clientless SSL VPN title bar appears above each web page. • You access websites by: – Entering the URL in the Enter Web Address field on the Clientless SSL VPN Home page. – Clicking on a preconfigured website link on the Clientless SSL VPN 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 web sites that appear as links on the Clientless SSL VPN Home page are available.

Table 38-2 Clientless SSL VPN 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 Clientless SSL VPN.
	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 38-2 Clientless SSL VPN Remote System Configuration and End User Requirements (continued)

Task	Remote System or End User Requirements	Specifications or Use Suggestions
Using Applications **(called Port Forwarding or Application Access)**	**Note** On Macintosh OS X, only the Safari browser supports this feature.	
	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.	
	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 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 Clientless SSL VPN 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 Clientless SSL VPN on the remote system and click the Application Access link on the Clientless SSL VPN 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 Clientless SSL VPN does not open the site over Clientless SSL VPN. To open a site over Clientless SSL VPN, cut and paste the URL into the Enter (URL) Address field.	

Table 38-2 Clientless SSL VPN 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 Clientless SSL VPN 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 Clientless SSL VPN.	
	Other mail clients	We have tested Microsoft Outlook Express versions 5.5 and 6.0. Clientless SSL VPN should support other SMTPS, POP3S, or IMAP4S e-mail programs via port forwarding, such as 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, 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 • Eudora 4.2 for Windows 2000 Other SSL-enabled mail clients should also work, but we have not verified them.
	Mail application configured	

Capturing Clientless SSL VPN Data

The CLI capture command lets you log information about websites that do not display properly over a Clientless SSL VPN 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 Clientless SSL VPN 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 Clientless SSL VPN session to a file.

Step 1 To start the Clientless SSL VPN 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 user logs in to begin a Clientless SSL VPN 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 Clientless SSL VPN 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 Clientless SSL VPN session and view it in a browser.

Step 1 To start the Clientless SSL VPN 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 user logs in to begin a Clientless SSL VPN 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 39

Clientless SSL VPN

Clientless SSL VPN 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. Clientless SSL VPN 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. Clientless SSL VPN 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 network resources on a user or group basis. Users have no direct access to these resources.

Clientless SSL VPN works on the platform in single, routed mode.

For information on configuring clientless SSL VPN for end users, see Clientless SSL VPN End User Set-up.

Security Precautions

Clientless SSL VPN connections on the security appliance differ from remote access IPSec connections, particularly with respect to how they interact with SSL-enabled servers, and precautions to follow to reduce security risks.

In a clientless SSL VPN connection, the security appliance acts as a proxy between the end user web browser and target web servers. When a user connects to an SSL-enabled web server, the security appliance establishes a secure connection and validates the server SSL certificate. The browser never receives the presented certificate, so it cannot examine and validate the certificate.

The current implementation of clientless SSL VPN on the security appliance does not permit communication with sites that present expired certificates. Nor does the security appliance perform trusted CA certificate validation to those SSL-enabled sites. Therefore, users do not benefit from certificate validation of pages delivered from an SSL-enabled web server before they use a web-enabled service.

⚠ **Caution** By default, the security appliance permits all portal traffic to all web resources (e.g., HTTPS, CIFS, RDP, and plug-ins). The security appliance clientless service rewrites each URL to one that is meaningful only to itself; the user cannot use the rewritten URL displayed on the page accessed to confirm that they are on the site they requested. To avoid placing users at risk, please assign a web ACL to the policies configured for clientless access – group-policies, dynamic access policies, or both – to control traffic flows from the portal. For example, without such an ACL, users could receive an authentication request from an outside fraudulent banking or commerce site. Also, we recommend disabling URL Entry on these policies to prevent user confusion over what is accessible. We recommend that you do the following to minimize risks posed by clientless SSL VPN access:

Step 1 Configure a group policy for all users who need clientless SSL VPN access, and enable clientless SSL VPN only for that group policy.

Step 2 With the group policy open, choose General > More Options > Web ACL and click **Manage**. Create a web ACL to do one of the following: permit access only to specific targets within the private network, permit access only to the private network, deny Internet access, or permit access only to reputable sites. Assign the web ACL to any policies (group policies, dynamic access policies, or both) that you have configured for clientless access. On a DAP, you select the web ACL on the Network ACL Filters tab.

Step 3 Disable URL entry on the *portal page*, the page that opens when the user establishes a browser-based connection. To do so, click Disable next to URL Entry on both the group policy Portal frame and the DAP Functions tab.

Step 4 Instruct users to enter external URLs in the native browser address field above the portal page or open a separate browser window to visit external sites.

Understanding Clientless SSL VPN System Requirements

The security appliance software version determines the endpoint OSs and browsers that clientless SSL VPN access supports. See Supported VPN Platforms, Cisco ASA 5500 Series for the list.

ActiveX pages require that you use the ActiveX Relay default setting (Enable) on the associated group policy. If you do so or assign a smart tunnel list to the policy, and the browser proxy exception list on the endpoint specifies a proxy, the user must add a "shutdown.webvpn.relay." entry to that list.

Browser-based VPN access does not support Windows Shares (CIFS) Web Folders on Windows 7, Vista, Internet Explorer 8, Mac OS, and Linux. Windows XP SP2 requires a Microsoft hotfix to support Web Folders.

See the following sections for the platforms supported by these clientless applications:

- Port Forwarding Requirements and Restrictions, page 39-22
- Smart Tunnel Requirements and Limitations, page 39-35
- Plug-in Requirements and Restrictions, page 39-68

Configuring Clientless SSL VPN Access

The Clientless SSL VPN Access pane lets you accomplish the following tasks:

- Enable or disable security appliance interfaces for clientless SSL VPN sessions.
- Choose a port for clientless SSL VPN connections.
- Set a global timeout value for clientless SSL VPN sessions.
- Set a maximum number of simultaneous clientless SSL VPN sessions.
- Configure the amount of security appliance memory that clientless SSL VPN can use.

To configure clientless SSL VPN services for individual users, the best practice is to choose the **Configuration > VPN > General > Group Policy >Add/Edit >WebVPN** pane. Then choose the **Configuration > Properties >Device Administration >User Accounts > VPN Policy** pane to assign the group policy to a user.

Fields

- Configure access parameters for WebVPN—Lets you enable or disable clientless SSL VPN connections on configured security appliance interfaces.
  - Interface—Displays names of all configured interfaces.
  - WebVPN Enabled—Displays current status for clientless SSL VPN on the interface.

    A green check next to Yes indicates that clientless SSL VPN is enabled.

    A red circle next to No indicates that clientless SSL VPN is disabled.
  - Enable/Disable—Click to enable or disable clientless SSL VPN on the highlighted interface.
- Port Number—Enter the port number that you want to use for clientless SSL VPN sessions. The default port is 443, for HTTPS traffic; the range is 1 through 65535. If you change the port number, All current clientless SSL VPN 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 clientless SSL VPN 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 clientless SSL VPN sessions you want to allow. Be aware that the different ASA models support clientless SSL VPN sessions as follows: ASA 5510 supports a maximum of 250; ASA 5520 maximum is 750; ASA 5540 maximum is 2500; ASA 5550 maximum is 5000.
- WebVPN Memory Size—Enter the percent of total memory or the amount of memory in kilobytes that you want to allocate to clientless SSL VPN 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, ASA 5550—4G. When you change the memory size, the new setting takes effect only after the system reboots.
- WebVPN Memory (unlabeled)—Choose to allocate memory for clientless SSL VPN either as a percentage of total memory or as an amount of memory in kilobytes.
- Enable Tunnel Group Drop-down List on WebVPN Login—Click to include a drop-down list of configured tunnel groups on the clientless SSL VPN 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

Clientless SSL VPN End User Set-up

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 ACE (access control entry), the security appliance denies it. ACEs are referred to as rules in this topic.

This pane lets you add and edit ACLs to be used for clientless SSL VPN sessions, and the ACL entries 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 ACLs to be applied to clientless SSL VPN sessions and their ACEs in the sequence determined by their position in the ACLs list 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.
- Enabled—Shows whether the ACE is enabled. When you create an ACE, by default it is enabled. Clear the check box to disable an ACE.
- 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 clientless SSL VPN 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.

For information about access rules and ACLs (including IPv6), see Information About Access Rules and ACLs, page 21-1.

For information about configuring access rules and ACLs (including IPv6), see Configuring Access Rules and ACLs, page 21-8.

For information about EtherType access rules and ACLs, see Chapter 21, "Configuring Access Rules and ACLs."

Fields

- ACL Name—Enter a name for the ACL. Maximum 55 characters.

Add/Edit ACE

An Access Control Entry (or "access rule") 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.

For information about access rules (including IPv6), see Information About Access Rules and ACLs, page 21-1.

For information about configuring access rules (including IPv6), see Configuring Access Rules and ACLs, page 21-8.

Fields

- Action—Permits or denies access to the specific networks, subnets, hosts, and web servers specified in the Filter group field.
- 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 field.
  - 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 field.
  - 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 field.
- 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 ACLs for clientless SSL VPN:

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
•	—	•	—	—

Configuring the Setup for Cisco Secure Desktop

The Cisco Secure Desktop Setup window displays the version and state of the Cisco Secure Desktop image if it is installed on the security appliance, indicates whether it is enabled, and shows the size of the cache used to hold the Cisco Secure Desktop and SSL VPN Client on the security appliance.

You can use the buttons in this window as follows:

- To transfer a copy of a Cisco Secure Desktop image from your local computer to the flash device of the security appliance, click **Upload**.

  To prepare to install or upgrade Cisco Secure Desktop, 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 your local computer to the flash device. Click **Browse Flash** to install it into the running configuration. Finally, click **Enable Secure Desktop**.

- To install or replace the Cisco Secure Desktop image on the flash device of the security appliance, click **Browse Flash**.

Note If you click **Browse Flash** to upgrade or downgrade the Cisco Secure Desktop image, select the package to install, and click **OK**, the Uninstall Cisco Secure Desktop dialog window asks you if you want to delete the Cisco Secure Desktop 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 Cisco Secure Desktop.

- To remove the Cisco Secure Desktop image and configuration file (sdesktop/data.xml) from the running configuration, click **Uninstall**.

  If you click this button, the Uninstall Cisco Secure Desktop dialog window asks if you want to delete the Cisco Secure Desktop image that was named in the "Secure Desktop Image field" and all Cisco Secure Desktop data files (including the entire Cisco Secure Desktop 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.

Fields

The Cisco Secure Desktop Setup pane displays the following fields:

- Location—Displays the Cisco Secure Desktop image loaded into the running configuration. By default, the filename is in the format securedesktop_asa_<n>_<n>*.pkg. Click **Browse Flash** to insert or modify the value in this field.
- Enable Secure Desktop—Click and click **Apply** to do the following:
  a. Make sure the file is a valid Cisco Secure Desktop image.
  b. Create an "sdesktop" folder on disk0 if one is not already present.
  c. Insert a data.xml (Cisco Secure Desktop 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 Cisco Secure Desktop to load the file.

  e. Enable Cisco Secure Desktop.

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

The Upload Image dialog box lets you transfer a copy of a Cisco Secure Desktop image from your local computer to the flash device on the security appliance. Use this window to install or upgrade Cisco Secure Desktop.

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.

You can use the buttons in this window as follows:

- To choose the path of the securedesktop_asa_<*n*>_<*n*>*.pkg file to be transferred, click **Browse Local Files**. 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**.
- To select the target directory for the file, click **Browse Flash**. The Browse Flash dialog box displays the contents of the flash card.
- To uploads the securedesktop_asa_<*n*>_<*n*>*.pkg file from your local computer to the flash device, click **Upload File**. 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 box removes the contents of the Local File Path and Flash File System Path fields.
- To close the Upload Image dialog box, click **Close**. Click this button after you upload the Cisco Secure Desktop image 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 Cisco Secure Desktop 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 Cisco Secure Desktop 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**.

Fields

The Upload Image dialog box displays the following fields:

- 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

  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
- File Name—Located in the Browse Flash dialog box that opens if you click **Browse Flash**, this field displays the name of the Cisco Secure Desktop image 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.

Modes

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

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

Configuring Application Helper

Clientless SSL VPN 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 Clientless SSL VPN 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.

Typically, Cisco TAC helps you write and apply an APCF.

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 pane 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—Click to rearrange APCF profiles within a list. The list determines the order in which the security appliance attempts to use APCF profiles.

Add/Edit APCF Profile

This pane 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—Click 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 pane 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 **Refresh**.

- 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—Click to use an APCF file stored on an HTTP, HTTPS or TFTP server.
- ftp, http, https, and tftp (unlabeled)—Identify the server type.
- URL (unlabeled)—Enter the path to the FTP, 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

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 Files—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, choose 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 **Close**.
- 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
•	—	•	—	—

Clock Accuracy for SharePoint Access

The clientless SSL VPN server on the security appliance uses cookies to interact with applications such as Microsoft Word on the endpoint. The cookie expiration time set by the security appliance can cause Word to malfunction when accessing documents on a SharePoint server if the time on the security appliance is incorrect. To prevent this malfunction, set the ASA clock properly. We recommend configuring the security appliance to dynamically synchronize with NTP services. For instructions, see "Setting the System Time."

Auto Signon

The Auto Signon window or tab lets you configure or edit auto signon for users of clientless SSL VPN. 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 user of clientless SSL VPN entered to log in 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 consists of authentication using Basic (HTTP), NTLM, FTP and CIFS, or all of these 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, or if you have Security Assertion Markup Language (SAML) Browser Post Profile SSO, and if you want to configure the security appliance to support this solution, see SSO Servers.

Note Do not enable auto signon for servers that do not require authentication or that use credentials different from the security appliance. When auto signon is enabled, the security appliance passes on the login credentials that the user entered to log into the security appliance regardless of what credentials are in user storage.

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, FTP and CIFS, or all of these methods—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

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 sign-on.
  - Mask—From the subnet mask menu, choose 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 Basic HTTP authentication requests, NTLM authentication requests, FTP and CIFS authentication requests, or requests using any of these methods.
  - Basic—Click this button if the servers support basic (HTTP) authentication.
  - NTLM—Click this button if the servers support NTLMv1 authentication.
  - FTP/CIFS—Click this button if the servers support FTP and CIFS authentication
  - Basic, NTLM, and FTP/CIFS—Click this button if the servers support all of the above.

Note If you configure one method for a range of servers (for example, HTTP Basic) and one of those servers attempts to authenticate with a different method (for example, NTLM), the security appliance does not pass the user 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
•	—	•	—	—

Configuring Session Settings

The Clientless SSL VPN Add/Edit Internal Group Policy > More Options > Session Settings window lets you specify personalized user information between clientless SSL VPN sessions. By default, each group policy inherits the settings from the default group policy. Use this window to specify personalized clientless SSL VPN user information for the default group policy and any group policies for which you want to differentiate these values.

Fields

- User Storage Location—Click none or choose the file server protocol (smb or ftp) from the drop-down menu. If you choose smb or ftp, use the following syntax to enter the file system destination into the adjacent text field:

  username:*password*@*host*:*port-number*/*path*

  For example

  mike:mysecret@ftpserver3:2323/public

  Note Although the configuration shows the username, password, and preshared key, the security appliance uses an internal algorithm to store the data in an encrypted form to safeguard it.

- Storage Key—Type the string, if required, for the security appliance to pass to provide user access to the storage location.
- Storage Objects—Choose one of the following options from the drop-down menu to specify the objects the server uses in association with the user. The security appliance store these objects to support clientless SSL VPN connections.
  - cookies,credentials
  - cookies
  - credentials
- Transaction Size-Enter the limit in KB over which to time out the session. This attribute applies only to a single transaction. Only a transaction larger than this value resets the session expiration clock.

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

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

Java Code Signer

Code signing appends a digital signature to the executable code itself. This digital signature provides enough information to authenticate the signer as well as to ensure that the code has not been subsequently modified since signed.

Code-signer certificates are special certificates whose associated private keys are used to create digital signatures. The certificates used to sign code are obtained from a CA, with the signed code itself revealing the certificate origin.

Choose a Java Code Signer from the drop down list.

To configure a Java Code Signer, choose **Configuration > Remote Access VPN > Certificate Management > Java Code Signer**.

Content Cache

Caching enhances the performance of clientless SSL VPN. It stores frequently reused objects in the system cache, which reduces the need to perform repeated rewriting and compressing of content. The use of the cache reduces traffic, with the result that many applications run more efficiently.

Fields

- Enable cache—Click to enable caching. The default value is disable.
- Parameters—Lets you define the terms for caching.
  - Enable caching of compressed content—Click 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.

  - 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.
  - 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.

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.

  - Cache static content—Click to cache all content that is not subject to rewrite, for example, PDF files and images.

- 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

The Content Rewrite pane lists all applications for which content rewrite is enabled or disabled.

Clientless SSL VPN 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.

By default, the security appliance rewrites, or transforms, all clientless traffic. 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.

"Example Content Rewrite Rules" shows example content rewrite rules.

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

- Enable content rewrite—Click to enable content rewrite for this rewrite rule.
- Rule Number—(Optional) Enter a number for this rule. This number specifies the priority of the rule, relative to the others 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 a string to match the application or resource to apply the rule to. The string can be up to 300 characters. You can use one of the following wildcards, but you must specify at least one alphanumeric character.

  * — Matches everything. ASDM does not accept a mask that consists of a * or *.*

  ? —Matches any single character.

  [!seq] — Matches any character not in sequence.

  [seq] — Matches any character in sequence.

Example Content Rewrite Rules

Function	Enable content rewrite	Rule Number	Rule Name	Resource Mask
Force all HTTP URLs to be delivered outside of ASA (split-tunneling)	Check	1	split-tunnel-all-http	http://*
Force all HTTPS URLs to be delivered outside of ASA	Check	2	split-tunnel-all-https	https://*

Modes

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

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

Java Code Signer

Java objects which have been transformed by clientless SSL VPN can subsequently be signed using a PKCS12 digital certificate associated with a trustpoint. In the Java Trustpoint pane, you can configure the clientless SSL VPN Java object signing facility to use a PKCS12 certificate and keying material from a specified trustpoint location. To import a trustpoint, choose **Configuration > Properties > Certificate > Trustpoint > Import**.

Fields

- Code Signer Certificate—Choose the configured certificate 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

This pane lets you view or specify the character encoding for clientless SSL VPN portal pages.

Character encoding, also called "character coding" and "a character set," is the pairing of raw data (such as 0s and 1s) with characters to represent the data. The language determines the character encoding method to use. Some languages use a single method, while others do not. Usually, the geographic region determines the default encoding method used by the browser, but the remote user can change it. 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 used on the portal page to ensure that the browser renders it properly, regardless of the region in which the user is using the browser, and regardless of any changes made to the browser.

By default, the security appliance applies the "Global Encoding Type" to pages from Common Internet File System servers. The mapping of CIFS servers to their appropriate character encoding, globally with the "Global 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, is an issue.

Fields

- Global Encoding Type —This attribute determines the character encoding that all clientless SSL VPN portal pages inherit except for those from the CIFS servers listed in the table. You can type the string or choose one of the options from the drop-down list, which contains 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 click **none** or specify a value that the browser on the clientless SSL VPN session 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 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 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

The Add CIFS Server Encoding dialog box lets you maintain exceptions to the “Global Encoding Type” attribute setting in the Add CIFS Encoding window. That pane 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 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 clientless SSL VPN portal pages. You can type the string, or choose 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 click **none** or specify a value that the browser on the clientless SSL VPN session 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
•	—	•	—	—

Web ACLs

The Web ACLs table displays the filters configured on the security appliance applicable to clientless SSL VPN traffic. The table shows the name of each access control list (ACL), and below and indented to the right of the ACL name, the ACEs (access control entries) assigned to the ACL.

Each ACL permits or denies access permits or denies access to specific networks, subnets, hosts, and web servers. Each ACE specifies one rule that serves the function of the ACL.

You can configure ACLs to apply to clientless SSL VPN traffic. The following rules apply:

- If you do not configure any filters, all connections are permitted.
- The security appliance supports only an inbound ACL on an interface.
- At the end of each ACL, an implicit, unwritten rule denies all traffic that is not explicitly permitted.

You can add ACLs and ACEs as follows:

- To add an ACL, click the down arrow next to the plus sign above the table and click **Add ACL**.

Note An ACL must be present before you can add an ACE.

- To add an ACE to an ACL that is already present in the table, choose it, then click the down arrow next to the plus sign above the table and click **Add ACE**.
- To insert an ACE before an ACE that is already present in the table, choose it, then click the down arrow next to the plus sign above the table and click **Insert**.
- To insert an ACE after an ACE that is already present in the table, choose it, then click the down arrow next to the plus sign above the table and click **Insert After**.

To change the values assigned to an ACE, double-click it, or choose it and click **Edit**.

To remove an ACL or an ACE, choose the entry in the table and click **Delete**.

The relative position of an ACE in an ACL determines the sequence with which the security appliance applies it to traffic on the interface. You can reorganize and reuse the ACEs present in the table as follows.

- To move an ACE above or below another ACE, choose it and click the up or down icon above the table.
- To move an ACE, choose the ACE, click the scissors icon above the table. Select the target ACL or ACE, click the arrow next to the clipboard icon, and click **Paste** to paste above the selection or **Paste After** to paste after the selection. The Edit ACE dialog box opens, providing you with an opportunity to change the values. Click **OK**.
- To copy an ACE, choose it and click the double-page icon above the table. Choose the target ACL or ACE, click the arrow next to the clipboard icon, and click **Paste** to paste above the selection or **Paste After** to paste after the selection. The Edit ACE dialog box opens, providing you with an opportunity to change the values. Click **OK**.

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

Both the Port Forwarding pane and Configure Port Forwarding Lists dialog box let you view the port forwarding lists. Both the Port Forwarding pane and the Add or Edit Port Forwarding Entry dialog box let you specify the name of a port forwarding list, and add, view, edit, and delete port forwarding entries to the list.

To add, change, or remove a port forwarding list, do one of the following:

- To add a port forwarding list and add entries to it, click **Add**. The Add Port Forwarding List dialog box opens. After you name the list, click **Add** again. ASDM opens the Add Port Forwarding Entry dialog box, which lets you assign the attributes of an entry to the list. After doing so and clicking OK, ASDM displays those attributes in the list. Repeat as needed to complete the list, then click **OK** in the Add Port Forwarding List dialog box.
- To change a port forwarding list, double-click the list or choose the list in the table and click **Edit**. Then click **Add** to insert a new entry into the list, or click an entry in the list and click **Edit** or **Delete**.
- To remove a list, select the list in the table and click **Delete**.

The following sections describe port forwarding and how to configure it:

- Why Port Forwarding?
- Port Forwarding Requirements and Restrictions
- Add/Edit Port Forwarding List
- Add/Edit Port Forwarding Entry

Why Port Forwarding?

Port forwarding is the legacy technology for supporting TCP-based applications over a clientless SSL VPN connection. You may choose to use port forwarding because you have built earlier configurations that support this technology.

Please consider the following alternatives to port forwarding:

- Smart tunnel access offers the following advantages to users:
  - Smart tunnel offers better performance than plug-ins.
  - Unlike port forwarding, smart tunnel simplifies the user experience by not requiring the user connection of the local application to the local port.
  - Unlike port forwarding, smart tunnel does not require users to have administrator privileges.
- Unlike port forwarding and smart tunnel access, a plug-in does not require the client application to be installed on the remote computer.

When configuring port forwarding on the security appliance, you specify the port the application uses. When configuring smart tunnel access, you specify the name of the executable file or its path.

Port Forwarding Requirements and Restrictions

In addition to the requirements in the "Understanding Clientless SSL VPN System Requirements" section on page 39-3, the following requirements and limitations apply to port forwarding on Windows:

- The remote host must be running a 32-bit version of one of the following:
  - Microsoft Windows Vista and Windows XP SP2 or SP3.
  - Apple Mac OS X 10.5 with Safari 2.0.4(419.3).
  - Fedora Core 4
- Browser-based users of Safari on Mac OS X 10.5.3 must identify a client certificate for use with the URL of the security appliance, once with the trailing slash and once without it, because of the way Safari interprets URLs. For example,
  - https://example.com/

  – https://example.com

  For details, go to the Safari, Mac OS X 10.5.3: Changes in client certificate authentication.

- Users of Microsoft Windows Vista who use port forwarding or smart tunnels must add the URL of the ASA to the Trusted Site zone. To access the Trusted Site zone, they must start Internet Explorer and choose the **Tools > Internet Options > Security** tab. Vista users can also disable Protected Mode to facilitate smart tunnel access; however, we recommend against this method because it increases the computer's vulnerability to attack.
- Port forwarding supports only 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 clientless SSL VPN port forwarding, but standard FTP, which uses ports 20 and 21, does not.
- Port forwarding does not support protocols that use UDP.
- The security appliance does not support the Microsoft Outlook Exchange (MAPI) proxy. For Microsoft Outlook Exchange communication using the MAPI protocol, remote users must use AnyConnect.
- A stateful failover does not retain sessions established using Application Access (either port forwarding or smart tunnel access). Users must reconnect following a failover.
- Port forwarding does not support connections to personal digital assistants.
- Port forwarding requires Sun JRE 5, Update 1.4 or later (JRE 6 or later recommended) to be enabled on the browser.

Caution 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 certificate store.

- 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.

  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.
- 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.

Configuring DNS for Port Forwarding

Port Forwarding forwards the domain name of the remote server or its IP address to the ASA for resolution and connection. In other words, the port forwarding applet accepts a request from the application and forwards it to the ASA. The ASA makes the appropriate DNS queries and establishes the connection on behalf of the port forwarding applet. The port forwarding applet only makes DNS queries to the ASA. It updates the host file so that when a port forwarding application attempts a DNS query, the query redirects to a loopback address.

Configure the security appliance to accept the DNS requests from the port forwarding applet as follows:

Step 1 Click **Configuration** > **Remote Access VPN** > **Clientless SSL VPN Access** > **Connection Profiles**.

The DefaultWEBVPNGroup entry is the default connection profile used for clientless connections.

Step 2 Highlight the DefaultWEBVPNGroup entry, then click **Edit** if your configuration uses it for clientless connections. Otherwise, highlight a connection profile used in your configuration for clientless connections, then click **Edit**.

The Basic window opens.

Step 3 Scan to the DNS area and select the DNS server from the drop-down list. Note the domain name, disregard the remaining steps, and go to the next section if ASDM displays the DNS server you want to use. You need to enter the same domain name when you specify the remote server while configuring an entry in the port forwarding list. Continue with the remaining steps if the DNS server is not present in the configuration.

Step 4 Click **Manage** in the DNS area.

The Configure DNS Server Groups window opens.

Step 5 Click **Configure Multiple DNS Server Groups**.

A window displays a table of DNS server entries.

Step 6 Click **Add.**

The Add DNS Server Group window opens.

Step 7 Enter a new server group name in the Name field, and enter the IP address and domain name (see Figure 39-1)

Figure 39-1 Example DNS Server Values for Port Forwarding

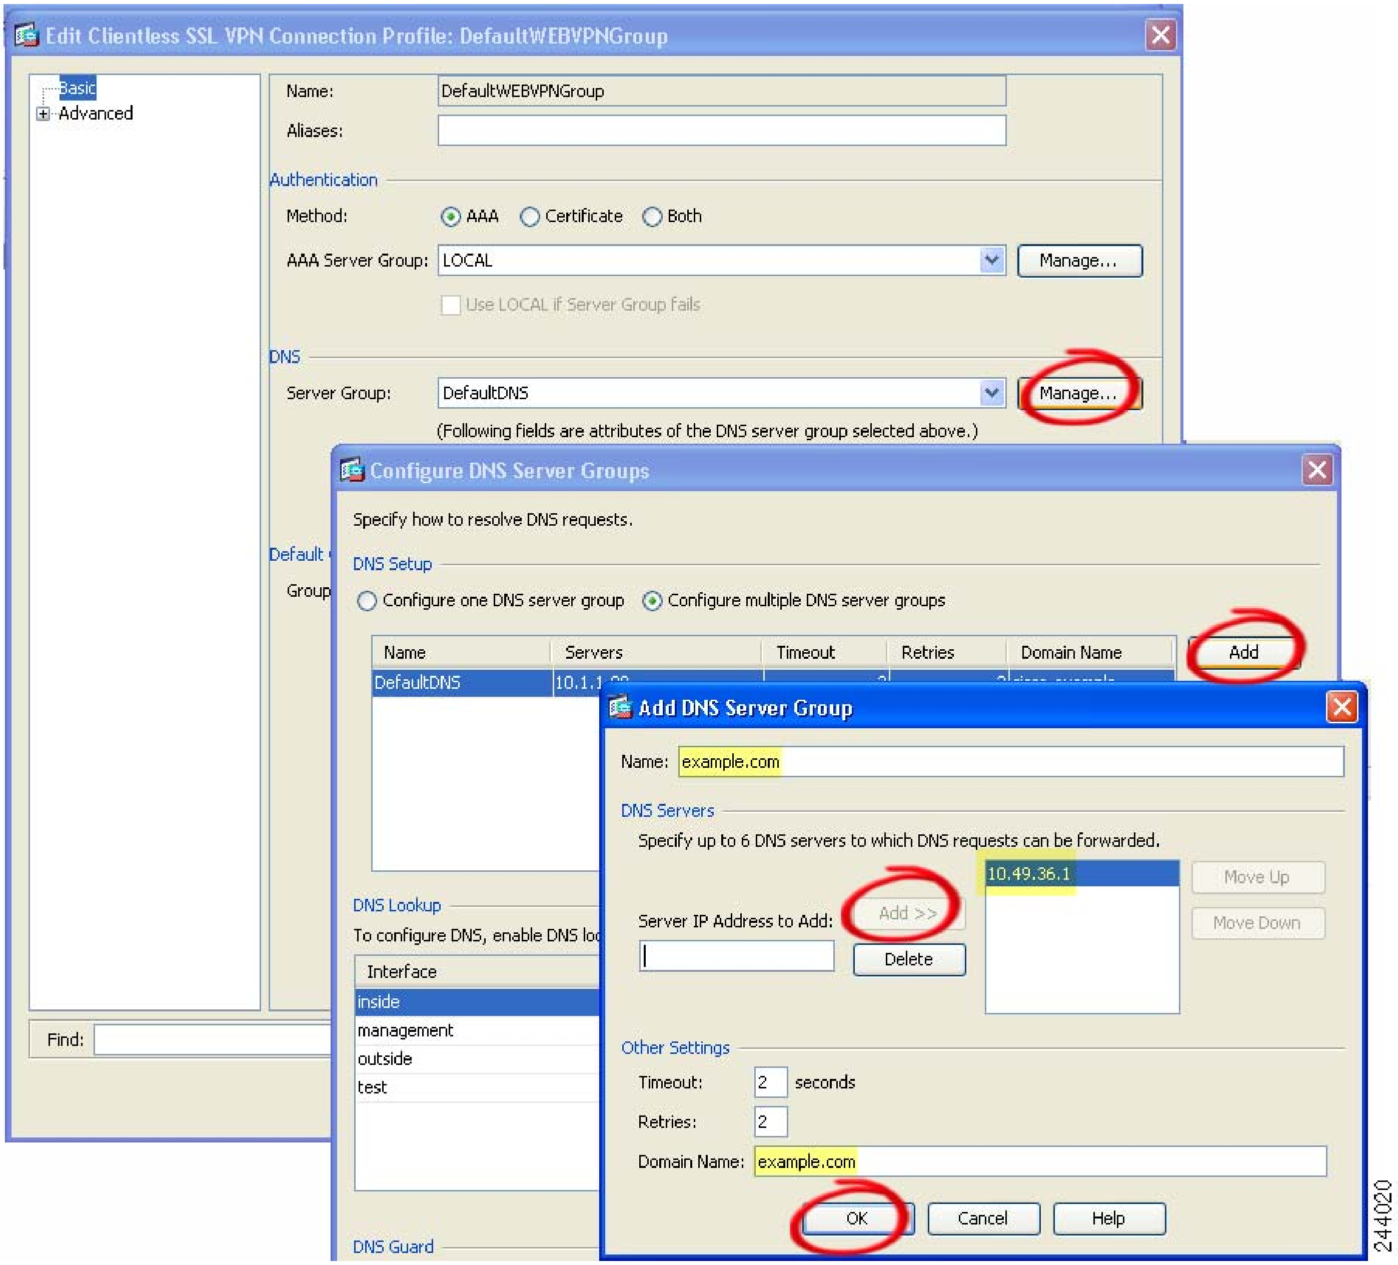

Edit Clientless SSL VPN Connection Profile: DefaultWEBVPNGroup

Name: DefaultWEBVPNGroup

Aliases:

Authentication

Method: ● AAA ○ Certificate ○ Both

AAA Server Group: LOCAL

□ Use LOCAL if Server Group fails

DNS

Server Group: DefaultDNS

(Following fields are attributes of the DNS server group selected above.)

Configure DNS Server Groups

Specify how to resolve DNS requests.

DNS Setup

○ Configure one DNS server group ● Configure multiple DNS server groups

Name	Servers	Timeout	Retries	Domain Name
DefaultDNS	10.1.1.99	2	2	cisco_example

DNS Lookup

To configure DNS, enable DNS loo

Interface
inside
management
outside
test

Add DNS Server Group

Name: example.com

DNS Servers

Specify up to 6 DNS servers to which DNS requests can be forwarded.

Server IP Address to Add:

10.49.36.1

Other Settings

Timeout: 2 seconds

Retries: 2

Domain Name: example.com

Note the domain name you entered. You need it when you specify the remote server later while configuring a port forwarding entry.

Step 8 Click **OK** until the Connection Profiles window becomes active again.

Step 9 Repeat Steps 2–8 for each remaining connection profile used in your configuration for clientless connections.

Step 10 Click **Apply**.

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

The Add/Edit Port Forwarding List dialog boxes let you add or edit a named list of TCP applications to associate with users or group policies for access over clientless SSL VPN connections.

Fields

- List Name—Alpha-numeric name for the list. Maximum 64 characters.
- Local TCP Port—Local port that listens for traffic for the application.
- Remote Server—IP address or DNS name of the remote server.
- Remote TCP Port—Remote port that listens for traffic for the application.
- Description—Text that describes the TCP application.

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

The Add/Edit Port Forwarding Entry dialog boxes let you specify TCP applications to associate with users or group policies for access over clientless SSL VPN connections. Assign values to the attributes in these windows as follows:

- Local TCP Port—Type a TCP 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—Enter either the domain name or IP address of the remote server. We recommend using a domain name so that you do not have to configure the client applications for the specific IP address.

Caution The DNS name assigned to the Remote Server parameter must match the Domain Name and Server Group parameters to establish the tunnel and resolve to an IP address, per the instructions in Add/Edit Port Forwarding List, page 39-26. The default setting for both the Domain and Server Group parameters is DefaultDNS.

- 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
•	—	•	—	—

Configuring the Use of External Proxy Servers

Use the Proxies pane to configure the security appliance to use external proxy servers to handle HTTP requests and HTTPS requests. These servers act as an intermediary between users and the Internet. Requiring all Internet access via servers you control provides another opportunity for filtering to assure secure Internet access and administrative control.

Note HTTP and HTTPS proxy services do not support connections to personal digital assistants.

Fields

Use an HTTP proxy server—Click to use an external HTTP proxy server.

- Specify IP address of proxy server—Click to identify the HTTP proxy server by its IP address or hostname.
- IP Address—Enter the hostname or IP address of the external HTTP proxy server.
- Port—Enter the port that listens for HTTP requests. The default port is 80.
- Exception Address List— (Optional) Enter a URL or a comma-delimited list of several URLs to exclude from those that can be sent to the HTTP proxy server. The string does not have a character limit, but the entire command cannot exceed 512 characters. You can specify literal URLs or use the following wildcards:
  - ***** to match any string, including slashes (/) and periods (.). You must accompany this wildcard with an alphanumeric string.
  - **?** to match any single character, including slashes and periods.
  - [*x*-*y*] to match any single character in the range of *x* and *y*, where *x* represents one character and *y* represents another character in the ANSI character set.
  - [**!***x*-*y*] to match any single character that is not in the range.

- UserName—(Optional) Enter this keyword to accompany each HTTP proxy request with a username to provide basic, proxy authentication.
- Password—Enter a password to send to the proxy server with each HTTP request.
- Specify PAC file URL—As an alternative to specifying the IP address of the HTTP proxy server, you can choose this option to specify a Proxy autoconfiguration file to download to the browser. Once downloaded, the PAC file uses a JavaScript function to identify a proxy for each URL. Enter **http://** and type the URL of the proxy autoconfiguration file into the adjacent field. If you omit the **http://** portion, the security appliance ignores it.

Use an HTTPS proxy server—Click to use an external HTTPS proxy server.

- Specify IP address of proxy server—Click to identify the HTTPS proxy server by its IP address or hostname.
- IP Address—Enter the hostname or IP address of the external HTTPS proxy server
- Port—Enter the port that listens for HTTPS requests. The default port is 443.
- Exception Address List— (Optional) Enter a URL or a comma-delimited list of several URLs to exclude from those that can be sent to the HTTPS proxy server. The string does not have a character limit, but the entire command cannot exceed 512 characters. You can specify literal URLs or use the following wildcards:
  - ***** to match any string, including slashes (/) and periods (.). You must accompany this wildcard with an alphanumeric string.
  - **?** to match any single character, including slashes and periods.
  - [*x*-*y*] to match any single character in the range of *x* and *y*, where *x* represents one character and *y* represents another character in the ANSI character set.
  - [**!***x*-*y*] to match any single character that is not in the range.
- UserName—(Optional) Enter this keyword to accompany each HTTPS proxy request with a username to provide basic, proxy authentication.
- Password—Enter a password to send to the proxy server with each HTTPS request.

Modes

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

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

Configuring 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.example.com/hrbenefits, *hrbenefits* is the path. Similarly, for the URL www.example.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

This pane 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—(radio button) Click to use a port for proxy bypass. The valid port numbers are 20000-21000.
  - Port (field)—Enter a high-numbered port for the security appliance to reserve for proxy bypass.
  - Path Mask—(radio button) Click to use a URL for proxy bypass.
  - Path Mask—(Field) Enter a URL for proxy bypass. It can contain a regular expression.
- URL—Define target URLs for proxy bypass.
  - URL—(drop-down list) Click either http or https as the protocol.
  - URL (text field)—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.

SSO Servers

The SSO Server pane lets you configure or delete single sign-on (SSO) for users of clientless SSL VPN connecting to a Computer Associates SiteMinder SSO server or to a Security Assertion Markup Language (SAML), Version 1.1, Browser Post Profile SSO server. SSO support, available only for clientless SSL VPN, lets users access different secure services on different servers without entering a username and password more than once.

You can choose from four methods when configuring SSO: Auto Signon using basic HTTP and/or NTLMv1 authentication, HTTP Form protocol, or Computer Associates eTrust SiteMinder (formerly Netegrity SiteMinder), or SAML, Version 1.1 Browser Post Profile.

Note The SAML Browser Artifact profile method of exchanging assertions is not supported.

The following sections describe the procedures for setting up SSO with both SiteMinder and SAML Browser Post Profile.

- Auto Signon—configures SSO with basic HTTP or NTLM authentication.
- Configuring Session Settings —configures SSO with the HTTP Form protocol.

The SSO mechanism either starts as part of the AAA process (HTTP Forms) or just after successful user authentication to either a AAA server (SiteMinder) or a SAML Browser Post Profile server. In these cases, the clientless SSL VPN server running on the security appliance acts as a proxy for the user to the authenticating server. When a user logs in, the clientless SSL VPN server sends an SSO authentication request, including username and password, to the authenticating server using HTTPS.

If the authenticating server approves the authentication request, it returns an SSO authentication cookie to the clientless SSL VPN 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.

Configuring SiteMinder and SAML Browser Post Profile

SSO authentication with SiteMinder or with SAML Browser Post Profile is separate from AAA and occurs after the AAA process completes. To set up SiteMinder SSO for a user or group, you must first configure a AAA server (RADIUS, LDAP and so forth). After the AAA server authenticates the user, the clientless SSL VPN server uses HTTPS to send an authentication request to the SiteMinder SSO server.

In addition to configuring the security appliance, for SiteMinder SSO, you also must configure your CA SiteMinder Policy Server with the Cisco authentication scheme. See Adding the Cisco Authentication Scheme to SiteMinder.

For SAML Browser Post Profile you must configure a Web Agent (Protected Resource URL) for authentication. For the specifics of setting up a SAML Browser Post Profile SSO server, see SAML POST SSO Server Configuration.

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 and the SAML Browser Post Profile 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
•	—	•	—	—

SAML POST SSO Server Configuration

Use the SAML server documentation provided by the server software vendor to configure the SAML server in Relying Party mode. To configure the SAML Server for Browser Post Profile, perform the following steps:

Step 1 Configure the SAML server parameters to represent the asserting party (the security appliance):

- Recipient consumer (Web Agent) URL (same as the assertion consumer URL configured on the ASA)
- Issuer ID, a string, usually the hostname of appliance
- Profile type -Browser Post Profile

Step 2 Configure certificates.

Step 3 Specify that asserting party assertions must be signed.

Step 4 Select how the SAML server identifies the user:

- Subject Name Type is DN

- Subject Name format is uid=<user>

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 steps:

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. This .jar file is also available on the Cisco security appliance CD.

Add/Edit SSO Servers

This SSO method uses CA SiteMinder and SAML Browser Post Profile. You can also set up SSO using the HTTP Form protocol, or Basic HTML and NTLM authentication. To use the HTTP Form protocol, see Configuring Session Settings. 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 types currently supported by the security appliance are SiteMinder and SAML Browser Post Profile.
- 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 the security appliance, the SSO server, 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
•	—	•	—	—

Configuring Smart Tunnel Access

The Smart Tunnels table displays the smart tunnel lists, each of which identifies one or more applications eligible for smart tunnel access, and its associated OS. Because each group policy or local user policy supports one smart tunnel list, you must group the nonbrowser-based applications to be supported into a smart tunnel list. Following the configuration of a list, you can assign it to one or more group polices or local user policies.

The Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Smart Tunnels window lets you do the following:

- To add a smart tunnel list and add applications to the list, click **Add**. The Add Smart Tunnel List dialog box opens. After you name the list, click **Add** again. ASDM opens the Add Smart Tunnel Entry dialog box, which lets you assign the attributes of a smart tunnel to the list. After doing so and clicking OK, ASDM displays those attributes in the list. Repeat as needed to complete the list, then click **OK** in the Add Smart Tunnel List dialog box.
- To change a smart tunnel list, double-click the list or choose the list in the table and click **Edit**. Then click **Add** to insert a new set of smart tunnel attributes into the list, or choose an entry in the list and click **Edit** or **Delete**.
- To remove a list, choose the list in the table and click **Delete**.

Following the configuration and assignment of a smart tunnel list, you can make a smart tunnel easy to use by adding a bookmark for the service and clicking the Enable Smart Tunnel Option in the Add or Edit Bookmark dialog box.

The following sections describe smart tunnel access and how to configure it:

- About Smart Tunnels
- Why Smart Tunnels?
- Smart Tunnel Requirements and Limitations

- Configuring a Smart Tunnel (Lotus example)
- Add or Edit Smart Tunnel List
- Add or Edit Smart Tunnel Entry
- Example Smart Tunnel Entries
- Add or Edit Smart Tunnel Auto Sign-on Server List
- Add or Edit Smart Tunnel Auto Sign-on Server Entry
- Example Smart Tunnel Entries

About Smart Tunnels

A smart tunnel is a connection between a TCP-based application and a private site, using a clientless (browser-based) SSL VPN session with the security appliance as the pathway, and the security appliance as a proxy server. You can identify applications to which you want to grant smart tunnel access, and specify the local path to each application. For applications running on Microsoft Windows, you can also require a match of the SHA-1 hash of the checksum as a condition for granting smart tunnel access.

Lotus SameTime and Microsoft Outlook Express are examples of applications to which you might want to grant smart tunnel access.

Configuring smart tunnels requires one of the following procedures, depending on whether the application is a client or is a web-enabled application:

- Create one or more smart tunnel lists of the client applications, then assign the list to the group policies or local user policies for whom you want to provide smart tunnel access.
- Create one or more bookmark list entries that specify the URLs of the web-enabled applications eligible for smart tunnel access, then assign the list to the DAPs, group policies, or local user policies for whom you want to provide smart tunnel access.

You can also list web-enabled applications for which to automate the submission of login credentials in smart tunnel connections over clientless SSL VPN sessions.

The following sections provide more information about smart tunnel access and how to configure it:

- Smart Tunnel Requirements and Limitations
- Configuring a Smart Tunnel (Lotus example)
- Add or Edit Smart Tunnel List
- Add or Edit Smart Tunnel Entry
- Example Smart Tunnel Entries
- Add or Edit Smart Tunnel Auto Sign-on Server List
- Add or Edit Smart Tunnel Auto Sign-on Server Entry

Why Smart Tunnels?

Smart tunnel access lets a client TCP-based application use a browser-based VPN connection to connect to a service. It offers the following advantages to users, compared to plug-ins and the legacy technology, port forwarding:

- Smart tunnel offers better performance than plug-ins.

- Unlike port forwarding, smart tunnel simplifies the user experience by not requiring the user connection of the local application to the local port.
- Unlike port forwarding, smart tunnel does not require users to have administrator privileges.

The advantage of a plug-in is that it does not require the client application to be installed on the remote computer.

Smart Tunnel Requirements and Limitations

The following sections categorize the smart tunnel requirements and limitations:

- General Requirements and Limitations
- Windows Requirements and Limitations
- Mac OS Requirements and Limitations

General Requirements and Limitations

Smart tunnel has the following general requirements and limitations:

- The remote host originating the smart tunnel must be running a 32-bit version of Microsoft Windows Vista, Windows XP, or Windows 2000; or Mac OS 10.4 or 10.5.
- Smart tunnel auto sign-on supports only Microsoft Internet Explorer on Windows.
- The browser must be enabled with Java, Microsoft ActiveX, or both.
- Smart tunnel supports only proxies placed between computers running Microsoft Windows and the security appliance. Smart tunnel uses the Internet Explorer configuration (that is, the one intended for system-wide use in Windows). If the remote computer requires a proxy server to reach the security appliance, the URL of the terminating end of the connection must be in the list of URLs excluded from proxy services. If the proxy configuration specifies that traffic destined for the ASA goes through a proxy, all smart tunnel traffic goes through the proxy.

  In an HTTP-based remote access scenario, sometimes a subnet does not provide user access to the VPN gateway. In this case, a proxy placed in front of the ASA to route traffic between the web and the end user's location provides web access. However, only VPN users can configure proxies placed in front of the ASA. When doing so, they must make sure these proxies support the CONNECT method. For proxies that require authentication, smart tunnel supports only the basic digest authentication type.
- When smart tunnel starts, the security appliance by default passes all browser traffic through the VPN session if the browser process is the same. The security appliance also does this if a tunnel-all policy applies. If the user starts another instance of the browser process, it passes all traffic through the VPN session. If the browser process is the same and the security appliance does not provide access to a URL, the user cannot open it. As a workaround, assign a tunnel policy that is not tunnel-all.
- A stateful failover does not retain smart tunnel connections. Users must reconnect following a failover.

Windows Requirements and Limitations

The following requirements and limitations apply to Windows only:

- Only Winsock 2, TCP-based applications are eligible for smart tunnel access.

- The security appliance does not support the Microsoft Outlook Exchange (MAPI) proxy. Neither port forwarding nor the smart tunnel supports MAPI. For Microsoft Outlook Exchange communication using the MAPI protocol, remote users must use AnyConnect.
- Users of Microsoft Windows Vista who use smart tunnel or port forwarding must add the URL of the ASA to the Trusted Site zone. To access the Trusted Site zone, they must start Internet Explorer and choose the **Tools** > **Internet Options** > **Security** tab. Vista users can also disable Protected Mode to facilitate smart tunnel access; however, we recommend against this method because it increases vulnerability to attack.

Mac OS Requirements and Limitations

The following requirements and limitations apply to Mac OS only:

- Safari 3.1.1 or later, or Firefox 3.0 or later.
- Sun JRE 1.5 or later.
- Only applications started from the portal page can establish smart tunnel connections. This requirement includes smart tunnel support for Firefox. Using Firefox to start another instance of Firefox during the first use of a smart tunnel requires the user profile named csco_st. If this user profile is not present, the session prompts the user to create one.
- Applications using TCP that are dynamically linked to the SSL library can work over a smart tunnel.
- The PowerPC MAC operating system is not supported with smart tunnel.
- Smart tunnel does not support the following on Mac OS:
  - Proxy services.
  - Auto sign-on.
  - Applications that use two-level name spaces.
  - Console-based applications, such as Telnet, SSH, and cURL.
  - Applications using dlopen or dlsym to locate libsocket calls.
  - Statically linked applications to locate libsocket calls.

Configuring a Smart Tunnel (Lotus example)

To configure a Smart Tunnel, perform the following steps:

Note These example instructions provide the minimum instructions required to add smart tunnel support for an application. See the field descriptions in the sections that follow for more information.

Step 1 Choose **Configuration** > **Remote Access VPN** > **Clientless SSL VPN Access** > **Portal** > **Smart Tunnels**.

Step 2 Double-click the smart tunnel list to which you want to add an application; or click **Add** to create a list of applications, enter a name for this list in the List Name field, and click **Add**.

For example, click **Add** in the Smart Tunnels pane, enter Lotus in the List Name field, and click **Add**.

Step 3 Click **Add** in the Add or Edit Smart Tunnel List dialog box.

Step 4 Enter a string in the Application ID field to serve as a unique index to the entry within the smart tunnel list.

Step 5 Enter the filename and extension of the application into the Process Name dialog box.

Table 39-1 shows example Application ID strings and the associated paths required to support Lotus.

Table 39-1 Smart Tunnel Example: Lotus 6.0 Thick Client with Domino Server 6.5.5

Application ID Example	Minimum Required Process Name
lotusnotes	notes.exe
lotusnlnotes	nlnotes.exe
lotusntaskldr	ntaskldr.exe
lotusnfileret	nfileret.exe

Step 6 Select **Windows** next to OS.

Step 7 Click **OK**.

Step 8 Repeat Steps 3–7 for each application to add to the list.

Step 9 Click **OK** in the Add or Edit Smart Tunnel List dialog box.

Step 10 Assign the list to the group policies and local user policies to which you want to provide smart tunnel access to the associated applications, as follows:

- To assign the list to a group policy, choose **Configuration > Remote Access VPN> Clientless SSL VPN Access > Group Policies > Add** or **Edit > Portal** and choose the smart tunnel name from the drop-down list next to the Smart Tunnel List attribute.
- To assign the list to a local user policy, choose **Configuration > Remote Access VPN> AAA Setup > Local Users > Add** or **Edit > VPN Policy > Clientless SSL VPN** and choose the smart tunnel name from the drop-down list next to the Smart Tunnel List attribute.

Add or Edit Smart Tunnel List

The Add Smart Tunnel List dialog box lets you add a list of smart tunnel entries to the security appliance configuration. The Edit Smart Tunnel List dialog box lets you modify the contents of the list.

Field

- List Name—Enter a unique name for the list of applications or programs. There is no restriction on the number of characters in the name. Do not use spaces.

  Following the configuration of the smart tunnel list, the list name appears next to the Smart Tunnel List attribute in the Clientless SSL VPN group policies and local user policies. Assign a name that will help you to distinguish its contents or purpose from other lists that you are likely to configure.

Modes

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

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

The following sections provide more information about smart tunnel access and how to configure it:

- Why Smart Tunnels?
- Smart Tunnel Requirements and Limitations
- Configuring a Smart Tunnel (Lotus example)
- Add or Edit Smart Tunnel Entry
- Example Smart Tunnel Entries
- Add or Edit Smart Tunnel Auto Sign-on Server List
- Add or Edit Smart Tunnel Auto Sign-on Server Entry

Add or Edit Smart Tunnel Entry

The Add or Edit Smart Tunnel Entry dialog box lets you specify the attributes of an application in a smart tunnel list.

- Application ID—Enter a string to name the entry in the smart tunnel list. The string is unique for the OS. It typically names the application to be granted smart tunnel access. To support multiple versions of an application for which you choose to specify different paths or hash values, you can use this attribute to differentiate entries, specifying the OS, and name and version of the application supported by each list entry. The string can be up to 64 characters.
- Process Name—Enter the filename or path to the application. The string can be up to 128 characters.

  Windows requires an exact match of this value to the right side of the application path on the remote host to qualify the application for smart tunnel access. If you specify only the filename for Windows, SSL VPN does not enforce a location restriction on the remote host to qualify the application for smart tunnel access.

  If you specify a path and the user installed the application in another location, that application does not qualify. The application can reside on any path as long as the right side of the string matches the value you enter.

  To authorize an application for smart tunnel access if it is present on one of several paths on the remote host, either specify only the name and extension of the application in this field; or create a unique smart tunnel entry for each path.

  Note A sudden problem with smart tunnel access may be an indication that a *Process Name* value is not up-to-date with an application upgrade. For example, the default path to an application sometimes changes following the acquisition of the company that produces the application and the next application upgrade.

For Windows, if you want to add smart tunnel access to an application started from the command prompt, you must specify “cmd.exe” in the Process Name of one entry in the smart tunnel list, and specify the path to the application itself in another entry, because “cmd.exe” is the parent of the application.

Mac OS requires the full path to the process, and is case-sensitive. To avoid specifying a path for each username, insert a tilde (~) before the partial path (e.g., ~/bin/vnc).

- OS—Click **Windows** or **Mac** to specify the host OS of the application.
- Hash—(Optional and applicable only for Windows) To obtain this value, enter the checksum of the application (that is, the checksum of the executable file) into a utility that calculates a hash using the SHA-1 algorithm. One example of such a utility is the Microsoft File Checksum Integrity Verifier (FCIV), which is available at http://support.microsoft.com/kb/841290/. After installing FCIV, place a temporary copy of the application to be hashed on a path that contains no spaces (for example, c:/fciv.exe), then enter **fciv.exe -sha1** *application* at the command line (for example, **fciv.exe -sha1 c:\msimn.exe**) to display the SHA-1 hash.

  The SHA-1 hash is always 40 hexadecimal characters.

  Before authorizing an application for smart tunnel access, clientless SSL VPN calculates the hash of the application matching the *Application ID*. It qualifies the application for smart tunnel access if the result matches the value of *Hash*.

  Entering a hash provides a reasonable assurance that SSL VPN does not qualify an illegitimate file that matches the string you specified in the *Application ID*. Because the checksum varies with each version or patch of an application, the *Hash* you enter can only match one version or patch on the remote host. To specify a hash for more than one version of an application, create a unique smart tunnel entry for each *Hash* value.

Note You must update the smart tunnel list in the future if you enter *Hash* values and you want to support future versions or patches of an application with smart tunnel access. A sudden problem with smart tunnel access may be an indication that the application list containing *Hash* values is not up-to-date with an application upgrade. You can avoid this problem by not entering a hash.

Following the configuration of the smart tunnel list, you must assign it to a group policy or a local user policy for it to become active, as follows:

- To assign the list to a group policy, choose **Config > Remote Access VPN> Clientless SSL VPN Access > Group Policies > Add** or **Edit > Portal** and choose the smart tunnel name from the drop-down list next to the Smart Tunnel List attribute.
- To assign the list to a local user policy, choose **Config > Remote Access VPN> AAA Setup > Local Users > Add** or **Edit > VPN Policy > Clientless SSL VPN** and choose the smart tunnel name from the drop-down list next to the Smart Tunnel List attribute.

Table 39-2 Example Smart Tunnel Entries

Smart Tunnel Support	Application ID (Any unique string is OK.)	Process Name	OS
Mozilla Firefox.	firefox	firefox.exe	Windows
Microsoft Outlook Express.	outlook-express	msimn.exe	Windows

Table 39-2 Example Smart Tunnel Entries

Smart Tunnel Support	Application ID (Any unique string is OK.)	Process Name	OS
More restrictive alternative—Microsoft Outlook Express only if the executable file is in a predefined path.	outlook-express	\Program Files\Outlook Express\msimn.exe	Windows
Open a new Terminal window on a Mac. (Any subsequent application launched from within the same Terminal window fails because of the one-time-password implementation.)	terminal	Terminal	Mac
Start smart tunnel for a new window	new-terminal	Terminal open -a MacTelnet	Mac
Start application from a Mac Terminal window.	curl	Terminal curl www.example.com	Mac

Modes

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

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

The following sections provide more information about smart tunnel access and how to configure it:

- Why Smart Tunnels?
- Smart Tunnel Requirements and Limitations
- Configuring a Smart Tunnel (Lotus example)
- Add or Edit Smart Tunnel List
- Add or Edit Smart Tunnel Auto Sign-on Server List
- Add or Edit Smart Tunnel Auto Sign-on Server Entry

Add or Edit Smart Tunnel Auto Sign-on Server List

The Add Smart Tunnel Auto Sign-on Server List dialog box lets you add one or more lists of servers for which to automate the submission of login credentials during smart tunnel setup. The Edit Smart Tunnel Auto-signon Server List dialog box lets you modify the contents of these lists.

Field

- List Name—Enter a unique name for the list of remote servers. The string can be up to 64 characters. Do not use spaces.

  Following the configuration of the smart tunnel auto sign-on list, the list name appears next to the Auto Sign-on Server List attribute under Smart Tunnel in the clientless SSL VPN group policy and local user policy configurations. Assign a name that will help you to distinguish its contents or purpose from other lists that you are likely to configure.

Modes

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

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

See also "Add or Edit Smart Tunnel Auto Sign-on Server Entry."

Add or Edit Smart Tunnel Auto Sign-on Server Entry

The Add or Edit Smart Tunnel Entry dialog box lets you identify a server to be added to a smart tunnel auto sign-on list. You can identify it by its hostname, or IP address and subnet mask.

Caution Use the address format used in the source code of the web pages on the intranet. If you are configuring smart tunnel auto sign-on for browser access and some web pages use host names and others use IP addresses, or you do not know, specify both in different smart tunnel auto sign-on entries. Otherwise, if a link on a web page uses a different format than the one you specify, it will fail when the user clicks it.

- Host name—Enter a hostname or wildcard mask to auto-authenticate to. You can use the following wildcard characters:
  - * to match any number of characters or zero characters
  - ? to match any single character
  - [] to match any single character in the range expressed inside the brackets

  For example, enter *.example.com. Using this option protects the configuration from dynamic changes to IP addresses.
- IP Address—Enter an IP address to auto-authenticate to.
- Subnet Mask—Sub-network of hosts associated with the IP address.
- Use Windows domain name with user name (Optional) —Click to add the Windows domain to the username if authentication requires it. If you do so, be sure to specify the domain name when assigning the smart tunnel list to one or more group policies or local user policies.

Following the configuration of the smart tunnel auto sign-on server list, you must assign it to a group policy or a local user policy for it to become active, as follows:

- To assign the list to a group policy, choose **Config > Remote Access VPN > Clientless SSL VPN Access > Group Policies > Add** or **Edit > Portal**, find the Smart Tunnel area, and choose the list name from the drop-down list next to the Auto Sign-on Server List attribute.
- To assign the list to a local user policy, choose **Config > Remote Access VPN> AAA Setup > Local Users > Add** or **Edit > VPN Policy > Clientless SSL VPN**, find the Smart Tunnel area, and choose the list name from the drop-down list next to the Auto Sign-on Server List attribute.

Modes

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

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

Configuring Customization Objects

You can customize all end-user visible content on the clientless SSL VPN portal. To do so, you create an XML customization object, using an XML template, the Customization Editor in ASDM, or by exporting and editing an already existing customization object, which you then reimport to the security appliance.

Version 8.0 software extends the functionality for configuring customization, and the new process is incompatible with previous versions. During the upgrade to 8.0 software, the security appliance preserves a current configuration by using old settings to generate new customization objects. This process occurs only once, and is more than a simple transformation from the old format to the new one because the old values are only a partial subset of the new ones.

Note Version 7.2 portal customizations and URL lists work in the Beta 8.0 configuration only if clientless SSL VPN (WebVPN) is enabled on the appropriate interface in the Version 7.2(x) configuration file *before* you upgrade to Version 8.0.

In the current pane, you can add a new customization object, based on a template, or you can modify an already-imported customization object.

Fields

Add—Click to invoke the Add Customization pane, which lets you make a copy of the default customization object and save it with a unique name. Then you can use the ASDM SSL VPN Customization Editor to modify it to suit your requirements.

Edit—Click to edit an existing, highlighted customization object. Doing so invokes the SSL VPN Customization Editor.

Delete—Click to delete a customization object.

Import—Click to import a customization object, which is an XML file. For information about creating such an XML file, click this link: Creating XML-Based Portal Customization Objects and URL Lists.

Export—Click to export an exiting, highlighted customization object. Doing so lets you edit the object, and then reimport it to this security appliance or to another one.

Customization Objects—Lists the existing customization objects on the security appliance.

OnScreen Keyboard—Specify when to display the OnScreen Keyboard to end users. This keyboard provides additional security by eliminating the need to enter keystrokes on a physical keyboard for passwords when users log on or otherwise authenticate.

Modes

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

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

Add Customization Object

To add a customization object, create a copy of and provide a unique name for the DfltCustomization object. Then you can modify or edit it to meet your requirements.

Field

Customization Object Name—Enter a name for the new customization object. Maximum 64 characters, no spaces.

Modes

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

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

Import/Export Customization Object

You can import or export already-existing customization objects. Import an object that you want to apply to end users. Export a customization object already resident on the security appliance for editing purposes, after which you can reimport it.

Fields

- Customization Object Name—Identify the customization object by name. Maximum 64 characters, no spaces.
- Select a file—Choose the method by which you want to import or export the customization file.
  - Local computer—Choose this method to import a file that resides on the local PC.
  - Path—Provide the path to the file.
  - Browse Local Files—Browse to the path for the file.
  - Flash file system—Choose this method to export a file that resides on the security appliance.
  - Path—Provide the path to the file.
  - Browse Flash—Browse to the path for the file.
  - Remote server—Choose this option to import a customization file that resides on a remote server accessible from the security appliance.
  - Path—Identify the method to access the file (ftp, http, or https), and provide the path to the file.

- Import/Export Now—Click to import or export the file.

Modes

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

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

Creating XML-Based Portal Customization Objects and URL Lists

This section includes the following topics:

- Understanding the XML Customization File Structure
- Customization Example
- Using the Customization Template

Understanding the XML Customization File Structure

Table 39-3 presents the file structure for an XML customization object.

Note An empty tag <param></param> in an XML customization file is the equivalent of a CLI command with a trivial value:

```
(hostname)# param value “”
```

Absence of a parameter/tag results in a default/inherited value, while presence results in setting the parameter/tag value even it is an empty string.

Table 39-3 XML-Based Customization File Structure

Tag	Type	Values	Preset value	Description
custom	node	—	—	Root tag
auth-page	node	—	—	Tag-container of authentication page configuration
window	node	—	—	Browser window
title-text	string	Arbitrary string	empty string	—
title-panel	node	—	—	The page top pane with a logo and a text
mode	text	enable|disable	disable	—
text	text	Arbitrary string	empty string	—
logo-url	text	Arbitrary URL	empty image URL	—

Table 39-3 XML-Based Customization File Structure (continued)

copyright-panel	node	—	—	The page bottom pane with a copyright information
mode	text	enable|disable	disable	—
text	text	Arbitrary URL	empty string	—
info-panel	node	—	—	The pane with a custom text and image
mode	string	enable|disable	disable	—
image-position	string	above|below	above	The image position, relative to text
image-url	string	Arbitrary URL	empty image	—
text	string	Arbitrary string	empty string	—
logon-form	node	—	—	The form with username, password, group prompt
title-text	string	Arbitrary string	Logon	—
message-text	string	Arbitrary string	empty string	—
username-prompt-text	string	Arbitrary string	Username	—
password-prompt-text	string	Arbitrary string	Password	—
internal-password-prompt-text	string	Arbitrary string	Internal Password	—
group-prompt-text	string	Arbitrary string	Group	—
submit-button-text	string	Arbitrary string	Logon	
logout-form	node	—	—	The form with a logout message and the buttons to login or close the window
title-text	string	Arbitrary string	Logout	—
message-text	string	Arbitrary string	Empty string	—
login-button-text	string	Arbitrary string	Login	
close-button-text	string	Arbitrary string	Close window	—
language-selector	node	—	—	The drop-down list to select a language
mode	string	enable|disable	disable	—
title	text	—	Language	The prompt text to select language

Table 39-3 XML-Based Customization File Structure (continued)

language	node (multiple)	—	—	—
code	string	—	—	—
text	string	—	—	—
portal	node	—	—	Tag-container of the portal page configuration
window	node	—	—	see authentication page description
title-text	string	Arbitrary string	Empty string	—
title-panel	node	—	—	see authentication page description
mode	string	enable|disable	Disable	—
text	string	Arbitrary string	Empty string	—
logo-url	string	Arbitrary URL	Empty image URL	—
navigation-panel	node	—	—	The pane on the left with application tabs
mode	string	enable|disable	enable	—
application	node (multiple)	—	N/A	The node changes defaults for the configured (by id) application
id	string	For stock application web-access file-access app-access net-access help For ins: Unique plug-in	N/A	—
tab-title	string	—	N/A	—

Table 39-3 XML-Based Customization File Structure (continued)

order	number	—	N/A	Value used to sort elements. The default element order values have step 1000, 2000, 3000, etc. For example, to insert an element between the first and second element, use a value 1001 – 1999.
url-list-title	string	—	N/A	If the application has bookmarks, the title for the panel with grouped bookmarks
mode	string	enable|disable	N/A	v
toolbar	node	—	—	—
mode	string	enable|disable	Enable	—
prompt-box-title	string	Arbitrary string	Address	Title for URL prompt list
browse-button-text	string	Arbitrary string	Browse	Browse button text
logout-prompt-text	string	Arbitrary string	Logout	—
column	node (multiple)	—	—	One column will be shown by default
width	string	—	N/A	—
order	number	—	N/A	Value used to sort elements.
url-lists	node	—	—	URL lists are considered to be default elements on the portal home page, if they are not explicitly disabled
mode	string	group | nogroup	group	Modes: group – elements grouped by application type i.e. Web Bookmarks, File Bookmarks) no-group – url-lists are shown in separate panes disable – do not show URL lists by default

Table 39-3 XML-Based Customization File Structure (continued)

panel	node (multiple)	—	—	Allows to configure extra panes
mode	string	enable|disable	—	Used to temporarily disable the panel without removing its configuration
title	string	—	—	—
type	string	—	—	Supported types: RSS IMAGE TEXT HTML
url	string	—	—	URL for RSS,IMAGE or HTML type paned
url-mode	string	—	—	Modes: mangle, no-mangle
text	string	—	—	Text for TEXT type panes
column	number	—	—	—

Customization Example

The following example illustrates the following customization options:

- Hides tab for the File access application
- Changes title and order of Web Access application
- Defines two columns on the home page
- Adds an RSS pane
- Adds three panes (text, image, and html) at the top of second pane

```
<custom name="Default">
  <auth-page>

    <window>
          <title-text l10n="yes">title WebVPN Logon</title>
    </window>

    <title-panel>
         <mode>enable</mode>
         <text l10n="yes">EXAMPLE WebVPN</text>
         <logo-url>http://www.example.com/images/EXAMPLE.gif</logo-url>
    </title-panel>

    <copyright>
         <mode>enable</mode>
         <text l10n="yes">(c)Copyright, EXAMPLE Inc., 2006</text>
    </copyright>

    <info-panel>
        <mode>enable</mode>
        <image-url>/+CSCOE+/custom/EXAMPLE.jpg</image-url>
      <text l10n="yes">
            <![CDATA[
            <div>
            <b>Welcome to WebVPN !.</b>
            </div>
            ]]>
      </text>
    </info-panel>

    <logon-form>
      <form>
          <title-text l10n="yes">title WebVPN Logon</title>
          <message-text l10n="yes">message WebVPN Logon</title>
          <username-prompt-text l10n="yes">Username</username-prompt-text>
          <password-prompt-text l10n="yes">Password</password-prompt-text>
          <internal-password-prompt-text l10n="yes">Domain
password</internal-password-prompt-text>
          <group-prompt-text l10n="yes">Group</group-prompt-text>
          <submit-button-text l10n="yes">Logon</submit-button-text>
      </form>
    </logon-form>

    <logout-form>
      <form>
          <title-text l10n="yes">title WebVPN Logon</title>
          <message-text l10n="yes">message WebVPN Logon</title>
          <login-button-text l10n="yes">Login</login-button-text>
          <close-button-text l10n="yes">Logon</close-button-text>
      </form>
    </logout-form>

    <language-slector>
      <language>
          <code l10n="yes">code1</code>
          <text l10n="yes">text1</text>
      </language>
      <language>
          <code l10n="yes">code2</code>
          <text l10n="yes">text2</text>
      </language>
    </language-slector>

  </auth-page>

  <portal>

    <window>
          <title-text l10n="yes">title WebVPN Logon</title>
    </window>

    <title-panel>
         <mode>enable</mode>
         <text l10n="yes">EXAMPLE WebVPN</text>
         <logo-url>http://www.example.com/logo.gif</logo-url>
    </title-panel>

    <navigation-panel>
         <mode>enable</mode>
    </navigation-panel>
    <application>
            <id>file-access</id>
            <mode>disable</mode>
    </application>
    <application>
            <id>web-access</id>
            <tab-title>EXAMPLE Intranet</tab-title>
            <order>3001</order>
    </application>

   <column>
            <order>2</order>
            <width>40%</width>
  <column>
   <column>
            <order>1</order>
            <width>60%</width>
  <column>

  <url-lists>
    <mode>no-group</mode>
  </url-lists>

  <pane>
   <id>rss_pane</id>
    <type>RSS</type>
   <url>rss.example.com?id=78</url>
  </pane>

  <pane>
    <id>text_pane</id>
    <type>TEXT</type>
   <url>rss.example.com?id=78</url>
    <column>1</column>
    <row>0</row>
    <text>Welcome to EXAMPLE  WebVPN Service</text>
  </pane>

  <pane>
    <type>IMAGE</type>
   <url>http://www.example.com/logo.gif</url>
    <column>1</column>
    <row>2</row>
   </pane>

  <pane>
    <type>HTML</type>
   <title>EXAMPLE news</title>
   <url>http://www.example.com/news.html</url>
    <column>1</column>
    <row>3</row>
   </pane>

  </portal>

</custom>
```

Using the Customization Template

A customization template, named *Template*, contains all currently employed tags with corresponding comments that describe how to use them. Use the **export** command to download the customization template from the security appliance, as follows:

```
hostname# export webvpn customization Template tftp://webserver/default.xml
hostname#
```

You cannot change or delete the file *Template*. When you export it as in this example, you are saving it to a new name, *default.xml*. After you make your changes to this file, using it to create a customization object that meets the needs of your organization, you import it to the security appliance, either as *default.xml* or another name of your choosing. For example:

```
hostname# import webvpn customization General tftp://webserver/custom.xml
hostname#
```

where you import an XML object called *custom.xml* and name it *General* on the security appliance.

The Customization Template

The customization template, named *Template*, follows:

```
<?xml version="1.0" encoding="UTF-8" ?>
- <!--

Copyright (c) 2008,2009 by Cisco Systems, Inc.
All rights reserved.

Note: all white spaces in tag values are significant and preserved.


Tag: custom
Description: Root customization tag

Tag: custom/languages
Description: Contains list of languages, recognized by ASA
Value: string containing comma-separated language codes. Each language code is
       a set dash-separated alphanumeric characters, started with
       alpha-character (for example: en, en-us, irokese8-language-us)
Default value: en-us

Tag: custom/default-language
Description: Language code that is selected when the client and the server
             were not able to negotiate the language automatically.
             For example the set of languages configured in the browser
             is "en,ja", and the list of languages, specified by
             'custom/languages' tag is "cn,fr", the default-language will be
             used.
Value: string, containing one of the language coded, specified in
       'custom/languages' tag above.
Default value: en-us

*********************************************************

Tag: custom/auth-page
Description: Contains authentication page settings

*********************************************************
Tag: custom/auth-page/window
Description: Contains settings of the authentication page browser window
```

```
Tag: custom/auth-page/window/title-text
Description: The title of the browser window of the authentication page
Value: arbitrary string
Default value: Browser's default value

*********************************************************

Tag: custom/auth-page/title-panel
Description: Contains settings for the title panel

Tag: custom/auth-page/title-panel/mode
Description: The title panel mode
Value: enable|disable
Default value: disable

Tag: custom/auth-page/title-panel/text
Description: The title panel text.
Value: arbitrary string
Default value: empty string

Tag: custom/auth-page/title-panel/logo-url
Description: The URL of the logo image (imported via "import webvpn webcontent")
Value: URL string
Default value: empty image URL

Tag: custom/auth-page/title-panel/background-color
Description: The background color of the title panel
Value: HTML color format, for example #FFFFFF
Default value: #FFFFFF

Tag: custom/auth-page/title-panel/font-color
Description: The background color of the title panel
Value: HTML color format, for example #FFFFFF
Default value: #000000

Tag: custom/auth-page/title-panel/font-weight
Description: The font weight
Value: CSS font size value, for example bold, bolder,lighter etc.
Default value: empty string

Tag: custom/auth-page/title-panel/font-size
Description: The font size
Value: CSS font size value, for example 10pt, 8px, x-large, smaller etc.
Default value: empty string



Tag: custom/auth-page/title-panel/gradient
Description: Specifies using the background color gradient
Value: yes|no
Default value:no

Tag: custom/auth-page/title-panel/style
Description: CSS style of the title panel
Value: CSS style string
Default value: empty string

*********************************************************

Tag: custom/auth-page/copyright-panel
Description: Contains the copyright panel settings
```

```
Tag: custom/auth-page/copyright-panel/mode
Description: The copyright panel mode
Value: enable|disable
Default value: disable

Tag: custom/auth-page/copyright-panel/text
Description: The copyright panel text
Value: arbitrary string
Default value: empty string

*********************************************************

Tag: custom/auth-page/info-panel
Description: Contains information panel settings

Tag: custom/auth-page/info-panel/mode
Description: The information panel mode
Value: enable|disable
Default value: disable

Tag: custom/auth-page/info-panel/image-position
Description: Position of the image, above or below the informational panel text
Values: above|below
Default value: above

Tag: custom/auth-page/info-panel/image-url
Description: URL of the information panel image (imported via "import webvpn webcontent")
Value: URL string
Default value: empty image URL

Tag: custom/auth-page/info-panel/text
Description: Text of the information panel
Text: arbitrary string
Default value: empty string

*********************************************************

Tag: custom/auth-page/logon-form
Description: Contains logon form settings

Tag: custom/auth-page/logon-form/title-text
Description: The logon form title text
Value: arbitrary string
Default value: "Logon"

Tag: custom/auth-page/logon-form/message-text
Description: The message inside of the logon form
Value: arbitrary string
Default value: empty string

Tag: custom/auth-page/logon-form/username-prompt-text
Description: The username prompt text
Value: arbitrary string
Default value: "Username"

Tag: custom/auth-page/logon-form/password-prompt-text
Description: The password prompt text
Value: arbitrary string
Default value: "Password"

Tag: custom/auth-page/logon-form/internal-password-prompt-text
Description: The internal password prompt text
Value: arbitrary string
Default value: "Internal Password"
```

```
Tag: custom/auth-page/logon-form/group-prompt-text
Description: The group selector prompt text
Value: arbitrary string
Default value: "Group"


Tag: custom/auth-page/logon-form/submit-button-text
Description: The submit button text
Value: arbitrary string
Default value: "Logon"

Tag: custom/auth-page/logon-form/internal-password-first
Description: Sets internal password first in the order
Value: yes|no
Default value: no


Tag: custom/auth-page/logon-form/title-font-color
Description: The font color of the logon form title
Value: HTML color format, for example #FFFFFF
Default value: #000000

Tag: custom/auth-page/logon-form/title-background-color
Description: The background color of the logon form title
Value: HTML color format, for example #FFFFFF
Default value: #000000


Tag: custom/auth-page/logon-form/font-color
Description: The font color of the logon form
Value: HTML color format, for example #FFFFFF
Default value: #000000

Tag: custom/auth-page/logon-form/background-color
Description: The background color of the logon form
Value: HTML color format, for example #FFFFFF
Default value: #000000


*********************************************************

Tag: custom/auth-page/logout-form
Description: Contains the logout form settings

Tag: custom/auth-page/logout-form/title-text
Description: The logout form title text
Value: arbitrary string
Default value: "Logout"

Tag: custom/auth-page/logout-form/message-text
Description: The logout form message text
Value: arbitrary string
Default value: Goodbye.
               For your own security, please:
               Clear the browser's cache
               Delete any downloaded files
               Close the browser's window

Tag: custom/auth-page/logout-form/login-button-text
Description: The text of the button sending the user to the logon page
Value: arbitrary string
Default value: "Logon"
```

```
*********************************************************

Tag: custom/auth-page/language-selector
Description: Contains the language selector settings

Tag: custom/auth-page/language-selector/mode
Description: The language selector mode
Value: enable|disable
Default value: disable

Tag: custom/auth-page/language-selector/title
Description: The language selector title
Value: arbitrary string
Default value: empty string

Tag: custom/auth-page/language-selector/language (multiple)
Description: Contains the language settings

Tag: custom/auth-page/language-selector/language/code
Description: The code of the language
Value (required): The language code string

Tag: custom/auth-page/language-selector/language/text
Description: The text of the language in the language selector drop-down box
Value (required): arbitrary string

*********************************************************

Tag: custom/portal
Description: Contains portal page settings

*********************************************************

Tag: custom/portal/window
Description: Contains the portal page browser window settings

Tag: custom/portal/window/title-text
Description: The title of the browser window of the portal page
Value: arbitrary string
Default value: Browser's default value

*********************************************************

Tag: custom/portal/title-panel
Description: Contains settings for the title panel

Tag: custom/portal/title-panel/mode
Description: The title panel mode
Value: enable|disable
Default value: disable

Tag: custom/portal/title-panel/text
Description: The title panel text.
Value: arbitrary string
Default value: empty string

Tag: custom/portal/title-panel/logo-url
Description: The URL of the logo image (imported via "import webvpn webcontent")
Value: URL string
Default value: empty image URL

Tag: custom/portal/title-panel/background-color
Description: The background color of the title panel
Value: HTML color format, for example #FFFFFF
```

```
Default value: #FFFFFF

Tag: custom/auth-pa/title-panel/font-color
Description: The background color of the title panel
Value: HTML color format, for example #FFFFFF
Default value: #000000

Tag: custom/portal/title-panel/font-weight
Description: The font weight
Value: CSS font size value, for example bold, bolder,lighter etc.
Default value: empty string

Tag: custom/portal/title-panel/font-size
Description: The font size
Value: CSS font size value, for example 10pt, 8px, x-large, smaller etc.
Default value: empty string

Tag: custom/portal/title-panel/gradient
Description: Specifies using the background color gradient
Value: yes|no
Default value:no

Tag: custom/portal/title-panel/style
Description: CSS style for title text
Value: CSS style string
Default value: empty string

*********************************************************

Tag: custom/portal/application (multiple)
Description: Contains the application setting

Tag: custom/portal/application/mode
Description: The application mode
Value: enable|disable
Default value: enable

Tag: custom/portal/application/id
Description: The application ID. Standard application ID's are: home, web-access,
file-access, app-access, network-access, help
Value: The application ID string
Default value: empty string

Tag: custom/portal/application/tab-title
Description: The application tab text in the navigation panel
Value: arbitrary string
Default value: empty string

Tag: custom/portal/application/order
Description: The order of the application's tab in the navigation panel. Applications with
lesser order go first.
Value: arbitrary number
Default value: 1000

Tag: custom/portal/application/url-list-title
Description: The title of the application's URL list pane (in group mode)
Value: arbitrary string
Default value: Tab tite value concatenated with "Bookmarks"

*********************************************************

Tag: custom/portal/navigation-panel
Description: Contains the navigation panel settings
```

```
Tag: custom/portal/navigation-panel/mode
Description: The navigation panel mode
Value: enable|disable
Default value: enable

*********************************************************

Tag: custom/portal/toolbar
Description: Contains the toolbar settings

Tag: custom/portal/toolbar/mode
Description: The toolbar mode
Value: enable|disable
Default value: enable

Tag: custom/portal/toolbar/prompt-box-title
Description: The universal prompt box title
Value: arbitrary string
Default value: "Address"

Tag: custom/portal/toolbar/browse-button-text
Description: The browse button text
Value: arbitrary string
Default value: "Browse"

Tag: custom/portal/toolbar/logout-prompt-text
Description: The logout prompt text
Value: arbitrary string
Default value: "Logout"

*********************************************************

Tag: custom/portal/column (multiple)
Description: Contains settings of the home page column(s)

Tag: custom/portal/column/order
Description: The order the column from left to right. Columns with lesser order values go
first
Value: arbitrary number
Default value: 0

Tag: custom/portal/column/width
Description: The home page column width
Value: percent
Default value: default value set by browser
Note: The actual width may be increased by browser to accommodate content


*********************************************************

Tag: custom/portal/url-lists
Description: Contains settings for URL lists on the home page

Tag: custom/portal/url-lists/mode
Description: Specifies how to display URL lists on the home page:
             group URL lists by application (group) or
             show individual URL lists (nogroup).
             URL lists fill out cells of the configured columns, which are not taken
             by custom panes.
             Use the attribute value "nodisplay" to not show URL lists on the home page.

Value: group|nogroup|nodisplay
Default value: group
```

```
*********************************************************

Tag: custom/portal/pane (multiple)
Description: Contains settings of the custom pane on the home page

Tag: custom/portal/pane/mode
Description: The mode of the pane
Value: enable|disable
Default value: disable

Tag: custom/portal/pane/title
Description: The title of the pane
Value: arbitrary string
Default value: empty string

Tag: custom/portal/pane/notitle
Description: Hides pane's title bar
Value: yes|no
Default value: no

Tag: custom/portal/pane/type
Description: The type of the pane. Supported types:
             TEXT - inline arbitrary text, may contain HTML tags;
             HTML - HTML content specified by URL shown in the individual iframe;
             IMAGE - image specified by URL
             RSS - RSS feed specified by URL
Value: TEXT|HTML|IMAGE|RSS
Default value: TEXT

Tag: custom/portal/pane/url
Description: The URL for panes with type  HTML,IMAGE or RSS
Value: URL string
Default value: empty string

Tag: custom/portal/pane/text
Description: The text value for panes with type TEXT
Value: arbitrary string
Default value:empty string

Tag: custom/portal/pane/column
Description: The column where the pane located.
Value: arbitrary number
Default value: 1

Tag: custom/portal/pane/row
Description: The row where the pane is located
Value: arbitrary number
Default value: 1

Tag: custom/portal/pane/height
Description: The height of the pane
Value: number of pixels
Default value: default value set by browser


*********************************************************

Tag: custom/portal/browse-network-title
Description: The title of the browse network link
Value: arbitrary string
Default value: Browse Entire Network


Tag: custom/portal/access-network-title
```

```
Description: The title of the link to start a network access session
Value: arbitrary string
Default value: Start AnyConnect

-->
- <custom>
- <localization>
<languages>en,ja,zh,ru,ua</languages>
<default-language>en</default-language>
</localization>
- <auth-page>
- <window>
- <title-text l10n="yes">
- <![CDATA[
WebVPN Service
]]>
</title-text>
</window>
- <language-selector>
<mode>disable</mode>
<title l10n="yes">Language:</title>
- <language>
<code>en</code>
<text>English</text>
</language>
- <language>
<code>zh</code>
<text>?? (Chinese)</text>
</language>
- <language>
<code>ja</code>
<text>?? (Japanese)</text>
</language>
- <language>
<code>ru</code>
<text>??????? (Russian)</text>
</language>
- <language>
<code>ua</code>
<text>?????????? (Ukrainian)</text>
</language>
</language-selector>
- <logon-form>
- <title-text l10n="yes">
- <![CDATA[
Login
]]>
</title-text>
- <title-background-color>
- <![CDATA[
#666666
]]>
</title-background-color>
- <title-font-color>
- <![CDATA[
#ffffff
]]>
</title-font-color>
- <message-text l10n="yes">
- <![CDATA[
Please enter your username and password.
]]>
</message-text>
- <username-prompt-text l10n="yes">
- <![CDATA[
USERNAME:
]]>
</username-prompt-text>
- <password-prompt-text l10n="yes">
- <![CDATA[
PASSWORD:
]]>
</password-prompt-text>
<internal-password-prompt-text l10n="yes" />
<internal-password-first>no</internal-password-first>
- <group-prompt-text l10n="yes">
- <![CDATA[
GROUP:
]]>
</group-prompt-text>
- <submit-button-text l10n="yes">
- <![CDATA[
Login
]]>
</submit-button-text>
- <title-font-color>
- <![CDATA[
#ffffff
]]>
</title-font-color>
- <title-background-color>
- <![CDATA[
#666666
]]>
</title-background-color>
<font-color>#000000</font-color>
<background-color>#ffffff</background-color>
</logon-form>
- <logout-form>
- <title-text l10n="yes">
- <![CDATA[
Logout
]]>
</title-text>
- <message-text l10n="yes">
- <![CDATA[
Goodbye.
]]>
</message-text>
</logout-form>
- <title-panel>
<mode>enable</mode>
- <text l10n="yes">
- <![CDATA[
WebVPN Service
]]>
</text>
<logo-url l10n="yes">/+CSCOU+/csco_logo.gif</logo-url>
<gradient>yes</gradient>
<style />
- <background-color>
- <![CDATA[
#ffffff
]]>
</background-color>
- <font-size>
- <![CDATA[
larger
]]>
</font-size>
- <font-color>
- <![CDATA[
#800000
]]>
</font-color>
- <font-weight>
- <![CDATA[
bold
]]>
</font-weight>
</title-panel>
- <info-panel>
<mode>disable</mode>
<image-url l10n="yes">/+CSCOU+/clear.gif</image-url>
<image-position>above</image-position>
<text l10n="yes" />
</info-panel>
- <copyright-panel>
<mode>disable</mode>
<text l10n="yes" />
</copyright-panel>
</auth-page>
- <portal>
- <title-panel>
<mode>enable</mode>
- <text l10n="yes">
- <![CDATA[
WebVPN Service
]]>
</text>
<logo-url l10n="yes">/+CSCOU+/csco_logo.gif</logo-url>
<gradient>yes</gradient>
<style />
- <background-color>
- <![CDATA[
#ffffff
]]>
</background-color>
- <font-size>
- <![CDATA[
larger
]]>
</font-size>
- <font-color>
- <![CDATA[
#800000
]]>
</font-color>
- <font-weight>
- <![CDATA[
bold
]]>
</font-weight>
</title-panel>
<browse-network-title l10n="yes">Browse Entire Network</browse-network-title>
<access-network-title l10n="yes">Start AnyConnect</access-network-title>
- <application>
<mode>enable</mode>
<id>home</id>
<tab-title l10n="yes">Home</tab-title>
<order>1</order>
</application>
```

```
- <application>
<mode>enable</mode>
<id>web-access</id>
- <tab-title l10n="yes">
- <![CDATA[
Web Applications
]]>
</tab-title>
- <url-list-title l10n="yes">
- <![CDATA[
Web Bookmarks
]]>
</url-list-title>
<order>2</order>
</application>
- <application>
<mode>enable</mode>
<id>file-access</id>
- <tab-title l10n="yes">
- <![CDATA[
Browse Networks
]]>
</tab-title>
- <url-list-title l10n="yes">
- <![CDATA[
File Folder Bookmarks
]]>
</url-list-title>
<order>3</order>
</application>
- <application>
<mode>enable</mode>
<id>app-access</id>
- <tab-title l10n="yes">
- <![CDATA[
Application Access
]]>
</tab-title>
<order>4</order>
</application>
- <application>
<mode>enable</mode>
<id>net-access</id>
<tab-title l10n="yes">AnyConnect</tab-title>
<order>4</order>
</application>
- <application>
<mode>enable</mode>
<id>help</id>
<tab-title l10n="yes">Help</tab-title>
<order>1000000</order>
</application>
- <toolbar>
<mode>enable</mode>
<logout-prompt-text l10n="yes">Logout</logout-prompt-text>
<prompt-box-title l10n="yes">Address</prompt-box-title>
<browse-button-text l10n="yes">Browse</browse-button-text>
</toolbar>
- <column>
<width>100%</width>
<order>1</order>
</column>
- <pane>
<type>TEXT</type>
```

```
<mode>disable</mode>
<title />
<text />
<notitle />
<column />
<row />
<height />
</pane>
- <pane>
<type>IMAGE</type>
<mode>disable</mode>
<title />
<url l10n="yes" />
<notitle />
<column />
<row />
<height />
</pane>
- <pane>
<type>HTML</type>
<mode>disable</mode>
<title />
<url l10n="yes" />
<notitle />
<column />
<row />
<height />
</pane>
- <pane>
<type>RSS</type>
<mode>disable</mode>
<title />
<url l10n="yes" />
<notitle />
<column />
<row />
<height />
</pane>
- <url-lists>
<mode>group</mode>
</url-lists>
</portal>
</custom>
```

Help Customization

The security appliance displays help content on the application panes during clientless sessions. Each clientless application pane displays its own help file content using a predetermined filename. For example, the help content displayed on the Application Access panel is from the file named app-access-hlp.inc. Table 39-4 shows the clientless application panels and predetermined filenames for the help content.

Table 39-4 Clientless Applications

Application Type	Panel	Filename
Standard	Application Access	app-access-hlp.inc
Standard	Browse Networks	file-access-hlp.inc
Standard	AnyConnect Client	net-access-hlp.inc

Table 39-4 Clientless Applications

Application Type	Panel	Filename
Standard	Web Access	web-access-hlp.inc
Plug-in	MetaFrame Access	ica-hlp.inc
Plug-in	Terminal Servers	rdp-hlp.inc
Plug-in	Telnet/SSH Servers	ssh,telnet-hlp.inc
Plug-in	VNC Connections	vnc-hlp.inc

You can customize the help files provided by Cisco or create help files in other languages. Then use the Import button to copy them to the flash memory of the security appliance for display during subsequent clientless sessions. You can also export previously imported help content files, customize them, and reimport them to flash memory.

The following sections describe how to customize or create help content visible on clientless sessions:

- Customizing a Help File Provided by Cisco
- Creating Help Files for Languages Not Provided by Cisco

Fields

Import—Click to launch the Import Application Help Content dialog, where you can import new help content to flash memory for display during clientless sessions.

Export—Click to retrieve previously imported help content selected from the table.

Delete—Click to delete previously imported help content selected from the table.

Language—Displays the abbreviation of the language rendered by the browser. This field is *not* used for file translation; it indicates the language used in the file. To identify the name of a language associated with an abbreviation in the table, display the list of languages rendered by your browser. For example, a dialog window displays the languages and associated language codes when you use one of the following procedures:

- Open Internet Explorer and choose **Tools** > **Internet Options** > **Languages** > **Add**.
- Open Mozilla Firefox and choose **Tools** > **Options** > **Advanced** > **General**, click **Choose** next to Languages, and click **Select a language to add**.

Filename—Displays the filename the help content file was imported as.

Modes

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

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

Customizing a Help File Provided by Cisco

To customize a help file provided by Cisco, you need to get a copy of the file from the flash memory card first. Get the copy and customize it by performing the following steps:

Step 1 Use your browser to establish a clientless session with the security appliance.

Step 2 Display the help file by appending the string in "URL of Help File in Flash Memory of the Security Appliance" in Table 39-5, to the address of the security appliance, substituting *language* as described below, then press **Enter**.

Table 39-5 Help Files Provided by Cisco for Clientless Applications

Application Type	Panel	URL of Help File in Flash Memory of the Security Appliance
Standard	Application Access	/+CSCOE+/help/*language*/app-access-hlp.inc
Standard	Browse Networks	/+CSCOE+/help/*language*/file-access-hlp.inc
Standard	AnyConnect Client	/+CSCOE+/help/*language*/net-access-hlp.inc
Standard	Web Access	/+CSCOE+/help/*language*/web-access-hlp.inc
Plug-in	Terminal Servers	/+CSCOE+/help/*language*/rdp-hlp.inc
Plug-in	Telnet/SSH Servers	/+CSCOE+/help/*language*/ssh,telnet-hlp.inc
Plug-in	VNC Connections	/+CSCOE+/help/*language*/vnc-hlp.inc

language is the abbreviation for the language rendered by the browser. It is *not* used for file translation; it indicates the language used in the file. For help files provided by Cisco in English, enter the abbreviation **en**.

The following example address displays the English version of the Terminal Servers help:

https://*address_of_security_appliance***/+CSCOE+/help/en/rdp-hlp.inc**

Step 3 Choose **File** > **Save (Page) As**.

Caution Do not change the contents of the File name box.

Step 4 Change the Save as type option to "Web Page, HTML only" and click **Save**.

Step 5 Use your preferred HTML editor to customize the file.

Note You can use most HTML tags, but do *not* use tags that define the document and its structure (for example, do not use <html>, <title>, <body>, <head>, <h1>, <h2>, etc. You can use character tags, such as the <b> tag, and the <p>, <ol>, <ul>, and <li> tags to structure content.

Step 6 Save the file as HTML only, using the original filename and extension.

Step 7 Make sure the filename matches the one in Table 39-5, and that it does not have an extra filename extension.

Return to ASDM and choose **Configuration** > **Remote Access VPN** > **Clientless SSL VPN Access** > **Portal** > **Help Customization** > **Import** to import the modified help file into flash memory.

Creating Help Files for Languages Not Provided by Cisco

Use standard HTML to create help files in other languages. We recommend creating a separate folder for each language you want to support.

Note You can use most HTML tags, but do *not* use tags that define the document and its structure (for example, do not use <html>, <title>, <body>, <head>, <h1>, <h2>, etc. You can use character tags, such as the <b> tag, and the <p>, <ol>, <ul>, and <li> tags to structure content.

Save the file as HTML only. Use the filename in the Filename column of Table 39-4.

Return to ASDM and choose **Configuration > Remote Access VPN > Clientless SSL VPN Access** > **Portal** > **Help Customization** > **Import** to import the new help file into flash memory.

Import/Export Application Help Content

Use the Import Application Help Content dialog box to import help files to flash memory for display on the portal pages during clientless sessions. Use the Export Application Help Content dialog box to retrieve previously imported help files for subsequent editing.

Fields

Language—For the Import Application Help Content dialog box only, this field specifies the language rendered by the browser. (This Language field is inactive in the Export Application Help Content dialog box.) This field is not used for file translation; it indicates the language used in the file. Click the dots next to the Language field, double-click the row containing the language used in the help file in the Browse Language Code dialog box, confirm the abbreviation in the Language Code field matches the abbreviation in the row, and click **OK**. If the language for which you want to provide help content is not present in the Browse Language Code dialog box, enter the abbreviation for the language you want into the Language Code field and click **OK**, or enter it into the Language text box to the left of the dots. To identify the abbreviation for the language of a help file to be imported if it is not present in the Browse Language Code dialog box, display the list of languages and abbreviations rendered by your browser. For example, a dialog box displays the languages and associated language codes when you use one of the following procedures:

- Open Internet Explorer and choose **Tools > Internet Options > Languages > Add**.
- Open Mozilla Firefox and choose **Tools > Options > Advanced > General**, click **Choose** next to Languages, and click **Select a language to add**.

File Name—If you are importing, choose the filename from the drop-down list for the new help content file. If you are exporting, this field is unavailable.

Select a File—Configure the parameters for the source file (if importing) or destination file (if exporting):

Local computer—Indicate if the source or destination file is on a local computer:

- Path—Identify the path of the source or destination file.
- Browse Local Files—Click to browse the local computer for the source or destination file.

Flash file system—Indicate if the source or destination file is located in flash memory on the security appliance:

- Path—Identify the path of the source or destination file in flash memory.
- Browse Flash—Click to browse the flash memory for the source or destination file.

Remote server—Indicate if the source or destination file is on a remote server:

- Path—Choose the file transfer (copy) method, either ftp, tftp, or http (for importing only), and specify the path.

Modes

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

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

Configuring Browser Access to Client-Server Plug-ins

The Client-Server Plug-in table displays the plug-ins the security appliance makes available to browsers in clientless SSL VPN sessions.

To add, change, or remove a plug-in, do one of the following:

- To add a plug-in, click **Import**. The Import Plug-ins dialog box opens.
- To remove a plug-in, choose it and click **Delete**.

The following sections describe the integration of browser plug-ins for clientless SSL VPN browser access:

- About Installing Browser Plug-ins
- Plug-in Requirements and Restrictions
- Preparing the Security Appliance for a Plug-in
- Installing Plug-ins Redistributed by Cisco
- Assembling and Installing Third-Party Plug-ins—Example: Citrix

About Installing Browser Plug-ins

A browser plug-in is a separate program that a web browser invokes to perform a dedicated function, such as connect a client to a server within the browser window. The security appliance lets you import plug-ins for download to remote browsers in clientless SSL VPN sessions. Of course, Cisco tests the plug-ins it redistributes, and in some cases, tests the connectivity of plug-ins we cannot redistribute. However, we do not recommend importing plug-ins that support streaming media at this time.

Note Per the GNU General Public License (GPL), Cisco redistributes plug-ins without having made any changes to them. Per the GPL, Cisco cannot directly enhance these plug-ins.

The security appliance does the following when you install a plug-in onto the flash device:

- (Cisco-distributed plug-ins only) Unpacks the jar file specified in the *URL*.
- Writes the file to the csco-config/97/plugin directory on the security appliance file system.

- Populates the drop-down menu next to the URL attributes in ASDM.
- Enables the plug-in for all future clientless SSL VPN sessions, and adds a main menu option and an option to the drop-down menu next to the Address field of the portal page.

  Table 39-6 shows the changes to the main menu and address field of the portal page when you add the plug-ins described in the following sections.

Table 39-6 Effects of Plug-ins on the Clientless SSL VPN Portal Page

Plug-in	Main Menu Option Added to Portal Page	Address Field Option Added to Portal Page
ica	Citrix Client	citrix://
rdp	Terminal Servers	rdp://
rdp2	Terminal Servers Vista	rdp2://
ssh,telnet	SSH	ssh://
	Telnet	telnet://
vnc	VNC Client	vnc://

Note A secondary security appliance obtains the plug-ins from the primary security appliance.

When the user in a clientless SSL VPN session clicks the associated menu option on the portal page, the portal page displays a window to the interface and displays a help pane. The user can select the protocol displayed in the drop-down menu and enter the URL in the Address field to establish a connection.

Note Some Java plug-ins may report a status of connected or online even when a session to the destination service is not set up. The open-source plug-in reports the status, not the security appliance.

Before installing the first plug-in, you must follow the instructions in the next section.

Plug-in Requirements and Restrictions

Clientless SSL VPN must be enabled on the security appliance to provide remote access to the plug-ins.

The minimum access rights required for remote use belong to the guest privilege mode.

A stateful failover does not retain sessions established using plug-ins. Users must reconnect following a failover.

Preparing the Security Appliance for a Plug-in

Before installing a plug-in, prepare the security appliance by performing the following steps:

Step 1 Make sure clientless SSL VPN (“webvpn”) is enabled on a security appliance interface.

Step 2 Install an SSL certificate onto the security appliance interface to which remote users use a fully-qualified domain name (FQDN) to connect.

Note Do not specify an IP address as the common name (CN) for the SSL certificate. The remote user attempts to use the FQDN to communicate with the security appliance. The remote PC must be able to use DNS or an entry in the System32\drivers\etc\hosts file to resolve the FQDN.

See the section that identifies the type of plug-in you want to provide for clientless SSL VPN access.

- Installing Plug-ins Redistributed by Cisco
- Assembling and Installing Third-Party Plug-ins—Example: Citrix

Installing Plug-ins Redistributed by Cisco

Cisco redistributes the following open-source, Java-based components to be accessed as plug-ins for web browsers in clientless SSL VPN sessions:

Table 39-7 Plug-ins Redistributed by Cisco

Cisco Download Link	Protocol	Description	Source of Redistributed Plug-in
rdp-plugin.090915.jar	RDP	Accesses Microsoft Terminal Services hosted by Windows Vista and Windows 2003 R2. Supports Remote Desktop ActiveX Control. **Note** We recommend using this plug-in that supports both RDP and RDP2. Only versions up to 5.2 of the RDP and RDP2 protocols are supported. Version 5.2 and later are not supported.	Cisco redistributes this plug-in without any changes to it per GNU General Public License. The original source of the redistributed plug-in is http://properjavardp.sourceforge.net/
rdp2-plugin.090211.jar	RDP2	Accesses Microsoft Terminal Services hosted by Windows Vista and Windows 2003 R2. Supports Remote Desktop ActiveX Control. **Note** This legacy plug-in supports only RDP2.	Cisco redistributes this plug-in without any changes to it per the GNU General Public License. The original source of the redistributed plug-in is http://properjavardp.sourceforge.net/
rdp-plugin.080506.jar	RDP	Accesses Microsoft Terminal Services hosted by Windows 2003 R1. Supports Remote Desktop ActiveX Control. **Note** This legacy plug-in supports only RDP.	Cisco redistributes this plug-in without any changes to it per the GNU General Public License. The source of the redistributed plug-in is http://properjavardp.sourceforge.net/

Table 39-7 Plug-ins Redistributed by Cisco

Cisco Download Link	Protocol	Description	Source of Redistributed Plug-in
ssh-plugin.080430.jar	SSH	The Secure Shell-Telnet plug-in lets the remote user establish a Secure Shell or Telnet connection to a remote computer. **Note** Because keyboard-interactive authentication is not supported by JavaSSH, it cannot be supported with SSH plugin. (Keyboard interactive is a generic authentication method used to implement different authentication mechanisms.	Cisco redistributes this plug-in without any changes to it per the GNU General Public License. The web site containing the source of the redistributed plug-in is http://javassh.org/
vnc-plugin.080130.jar	VNC	The Virtual Network Computing plug-in lets the remote user use a monitor, keyboard, and mouse to view and control a computer with remote desktop sharing turned on. This version changes the default color of the text, and contains updated French and Japanese help files.	Cisco redistributes this plug-in without any changes to it per the GNU General Public License. The web site containing the source of the redistributed plug-in is http://www.tightvnc.com/

To retrieve a plug-in redistributed by Cisco and import it into the security appliance, perform the following steps:

Step 1 Create a temporary directory named plugins on the computer you use to establish ASDM sessions with the security appliance.

Step 2 Download the plug-ins you want from the Cisco website to the plugins directory.

Step 3 Choose **Configuration** > **Remote Access VPN** > **Clientless SSL VPN Access** > **Portal** > **Client-Server Plug-ins**.

This pane displays the plug-ins that are available to clientless SSL sessions.

Step 4 Click **Import**.

The Import Client-Server Plug-in dialog box opens.

Step 5 Use the following descriptions to enter the field values.

Fields

The Import Client-Server Plug-in dialog box displays the following fields:

- Plug-in Name—Select one of the following values:
  - **ica** to provide plug-in access to Citrix MetaFrame or Web Interface services. Then specify the path to the ica-plugin.jar file in the Remote Server field, as described below.
  - **rdp** to provide plug-in access to Remote Desktop Protocol services. Then specify the path to the rdp-plugin.jar file in the Remote Server field.
  - **ssh,telnet** to provide plug-in access to *both* Secure Shell and Telnet services. Then specify the path to the ssh-plugin.jar file in the Remote Server field.
  - **vnc** to provide plug-in access to Virtual Network Computing services. Then specify the path to the vnc-plugin.jar file in the Remote Server field.

Note Any undocumented options in this menu are experimental and are not supported.

- Select a file—Click one of the following options and insert a path into its text field.
  - Local computer—Click to retrieve the plug-in from the computer with which you have established the ASDM session. Enter the location and name of the plug-in into the associated Path field, or click **Browse Local Files** and navigate to the plug-in, choose it, then click **Select**.
  - Flash file system—Click if the plug-in is present on the file system of the security appliance. Enter the location and name of the plug-in into the associated Path field, or click **Browse Flash** and navigate to the plug-in, choose it, then click **OK**.
  - Remote Server—Click to retrieve the plug-in from a host running an FTP or TFTP server. Choose **ftp**, **tftp**, or **HTTP** from the drop-down menu next to the associated Path attribute, depending on which service is running on the remote server. Enter the host name or address of the server and the path to the plug-in into the adjacent text field.

Step 6 Click **Import Now**.

Click **Apply**.

The plug-in is now available for future clientless SSL VPN sessions.

Modes

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

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

Assembling and Installing Third-Party Plug-ins—Example: Citrix

The open framework of the security appliance lets you add plug-ins to support third-party Java client/server applications. As an example of how to provide clientless SSL VPN browser access to third-party plug-ins, this section describes how to add clientless SSL VPN support for the Citrix Presentation Server Client or Citrix Web Interface (for XenDesktop).

Caution Cisco does not provide direct support for or recommend any particular plug-ins that are not redistributed by Cisco. As a provider of clientless SSL VPN services, you are responsible for reviewing and complying with any license agreements required for the use of plug-ins.

With a Citrix plug-in installed on the security appliance, clientless SSL VPN users can use a connection to the security appliance to access Citrix MetaFrame or Web Interface services.

A stateful failover does not retain sessions established using the Citrix plug-in. Citrix users must reauthenticate after failover.

To provide access to the Citrix plug-in, follow the procedures in the following sections.

Preparing the Citrix MetraFrame Server for Clientless SSL VPN Access

The security appliance performs the connectivity functions of the Citrix secure gateway when the Citrix client connects to the Citrix MetaFrame Server or Web Interface. Therefore, you must configure the Citrix Web Interface software to operate in a mode that does not use the (Citrix) "secure gateway." Otherwise, the Citrix client cannot connect to the Citrix MetaFrame Server.

Follow the instructions in the "Preparing the Security Appliance for a Plug-in" section on page 39-68 before using the next section, if you are not already providing support for a plug-in.

Follow Steps 1 – 4 of http://support.citrix.com/article/CTX117597 if you are configuring access to Web Interface (for XenDesktop), or you later upgrade to it, to avoid Cookies Required errors.

Creating, Installing, and Testing the Citrix Plug-in

To create and install the Citrix plug-in, perform the following steps:

Step 1 Download the ica-plugin.zip file from the Cisco Software Download website.

This file contains files that Cisco customized for use with the Citrix plug-in.

Step 2 Download the Citrix Java client from the Citrix site.

Step 3 Extract the following files from the Citrix Java client:

- JICA-configN.jar
- JICAEngN.jar

You can use WinZip to perform this step and the next.

Step 4 Add the extracted files to the ica-plugin.zip file.

Step 5 Ensure the EULA included with the Citrix Java client grants you the rights and permissions to deploy the client on your web servers.

Step 6 Establish an ASDM session with the security appliance, choose **Config** > **Remote Access VPN** > **Clientless SSL VPN Access** > **Portal** > **Client-Server Plug-ins** > **Import**, and import the ica-plugin.zip file.

Note Users of clientless SSL VPN sessions cannot enter a URL in the Address box to get SSO support for Citrix sessions. You must insert a bookmark as instructed in the following step if you want to provide SSO support for the Citrix plug-in.

Step 7 Add a bookmark to the applicable bookmark list to make it easy for users to connect. Choose **ica** and enter the following information into the Address field:

citrix-server/**?DesiredColor=4&DesiredHRes=1024&DesiredVRes=768**

See Add/Edit Bookmark List and Add Bookmark Entry as needed.

Step 8 To test the plug-in, establish a clientless session with the security appliance and click the bookmark.

Use the Client for Java Administrator's Guide as needed.

Language Localization

The security appliance provides language translation for the portal and screens displayed to users that initiate browser-based, clientless SSL VPN connections, screens associated with optional plug-ins, and the interface displayed to Cisco AnyConnect VPN Client users.

This section describes how to configure the security appliance to translate these user messages and includes the following sections:

- Understanding Language Translation, page 39-73
- Creating a Translation Table, page 39-74
- Add/Edit Localization Entry, page 39-75
- Import/Export Language Localization, page 39-79

Understanding Language Translation

Each functional area and its messages that are visible to remote users are organized into translation domains. Table 39-8 shows the translation domains and the functional areas translated.

Table 39-8 Translation Domains and Functional Areas Affected

Translation Domain	Functional Areas Translated
AnyConnect	Messages displayed on the user interface of the Cisco AnyConnect VPN Client.
CSD	Messages for the Cisco Secure Desktop (CSD).
customization	Messages on the logon and logout pages, portal page, and all the messages customizable by the user.
keepout	Message displayed to remote users when VPN access is denied.
PortForwarder	Messages displayed to Port Forwarding users.
url-list	Text that user specifies for URL bookmarks on the portal page.
webvpn	All the layer 7, AAA and portal messages that are not customizable.
plugin-ica	Messages for the Citrix plug-in.
plugin-rdp	Messages for the Remote Desktop Protocol plug-in.
plugin-telnet,ssh	Messages for the Telnet and SSH plug-in.
plugin-vnc	Messages for the VNC plug-in.

The software image package for the security appliance includes a language localization template for each domain that is part of the standard functionality. The templates for plug-ins are included with the plug-ins and define their own translation domains.

You can export the template for a translation domain, which creates an XML file of the template at the URL you provide. The message fields are empty in this file. You can customize the messages and import the template to create a new language localization table that resides in flash memory.

You can also export an existing language localization table. The XML file created displays the messages you edited previously. Reimporting this XML file with the same language name creates a new version of the language localization table, overwriting previous messages.

Some templates are static, but some change based on the configuration of the security appliance. Because you can customize the *logon and logout pages, portal page, and URL bookmarks for clientless sessions,* the **security appliance generates the customization** and **url-list** translation domain templates dynamically and the template automatically reflects your changes to these functional areas.

After creating language localization tables, they are available to customization objects that you create and apply to group policies or user attributes. A language localization table has no affect and messages are not translated on user screens until you create the customization object, identify a language localization table to use in that object, and specify the customization for the group policy or user.

Fields

Add—Launches the Add Localization Entry dialog where you can select a localization template to add and you can edit the contents of the template.

Edit—Launches the Edit Localization Entry dialog for the selected language in the table, and allows you to edit the previously-imported language localization table.

Delete—Deletes a selected language localization table.

Import—Launches the Import Language Localization dialog where you can import a language localization template or table.

Export—Launches the Export Language Localization dialog where you can export a language localization template or table to a URL where you can make changes to the table or template.

Language—The language of existing Language Localization tables.

Language Localization Template—The template that the table is based on.

Creating a Translation Table

To create a translation table, perform the following steps:

Step 1 Choose **Remove Access VPN > Clientless SSL VPN Access > Portal > Advanced > Language Localization**. The Language Localization pane displays. Click **Add**. The Add Language Localization window displays.

Step 2 Choose a Language Localization Template from the drop-down box. The entries in the box correspond to functional areas that are translated. For more information about the functionality for each template, see table Table 39-7.

Step 3 Specify a language for the template. The template becomes a translation table in cache memory with the name you specify. Use an abbreviation that is compatible with the language options for your browser. For example, if you are creating a table for the Chinese language, and you are using IE, use the abbreviation *zh*, that is recognized by IE.

Step 4 Edit the translation table. For each message represented by the msgid field that you want to translate, enter the translated text between the quotes of the associated msgstr field. The example below shows the message Connected, with the Spanish text in the msgstr field:

```
msgid "Connected"
msgstr "Conectado"
```

Step 5 Click **OK**. The new table appears in the list of translation tables.

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 Localization Entry

You can add a new translation table, based on a template, or you can modify an already-imported translation table in this pane.

Fields

Language Localization Template—Select a template to modify and use as a basis for a new translation table. The templates are organized into translation domains and affect certain areas of functionality. The following table shows the translation domains and the functional areas affected:

Translation Domain	Functional Areas Translated
AnyConnect	Messages displayed on the user interface of the Cisco AnyConnect VPN client.
CSD	Messages for the Cisco Secure Desktop (CSD).
customization	Messages on the logon and logout pages, portal page, and all the messages customizable by the user.
keepout	Message displayed to remote users when VPN access is denied.
PortForwarder	Messages displayed to Port Forwarding users.
url-list	Text that user specifies for URL bookmarks on the portal page.
webvpn	All the layer 7, AAA and portal messages that are not customizable.
plugin-ica	Messages for the Citrix plug-in.
plugin-rdp	Messages for the Remote Desktop Protocol plug-in.
plugin-telnet,ssh	Messages for the Telnet and SSH plug-in.
plugin-vnc	Messages for the VNC plug-in.

Language—Specify a language. Use an abbreviation that is compatible with the language options of your browser. The security appliance creates the new translation table with this name.

Text Editor—Use the editor to change the message translations. The message ID field (msgid) contains the default translation. The message string field (msgstr) that follows msgid provides the translation. To create a translation, enter the translated text between the quotes of the msgstr string. For example, to translate the message “Connected” with a Spanish translation, insert the Spanish text between the msgstr quotes:

```
msgid "Connected"
msgstr "Conectado"
```

After making changes, click **Apply** to import the translation table.

Modes

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

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

AnyConnect Customization

Resources

Specify resource files that customize or re-brand the AnyConnect VPN client in this pane.

Note The security appliance does not support this feature for the AnyConnect VPN client, Versions 2.0 and 2.1. For more information on manually customizing the client, see the *AnyConnect VPN Client Administrator Guide* and the *Release Notes for Cisco AnyConnect VPN Client*.

Fields

Import—Launches the Import AnyConnect Customization Objects dialog, where you can specify a file to import as an object.

Export—Launches the Export AnyConnect Customization Objects dialog, where you can specify a file to export as an object.

Delete—Removes the selected object.

Platform—The type of remote PC platform supported by the object.

Object Name—The name of the object.

Modes

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

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

Binary

Specify third-party programs that use the AnyConnect VPN client API in this pane. The security appliance downloads these programs to the client for customizing the user interface or the command line interface.

Note The security appliance does not support this feature for the AnyConnect VPN client, Versions 2.0 and 2.1. For more information on manually customizing the client, see the *AnyConnect VPN Client Administrator Guide* and the *Release Notes for Cisco AnyConnect VPN Client*.

Fields

Import—Launches the Import AnyConnect Customization Objects dialog, where you can specify a file to import as an object.

Export—Launches the Export AnyConnect Customization Objects dialog, where you can specify a file to export as an object.

Delete—Removes the selected object.

Platform—The type of remote PC platform supported by the object.

Object Name—The name of the object.

Modes

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

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

GUI Text and Messages

Change text and messages displayed on the AnyConnect client GUI displayed to remote users in this pane. This pane also shares functionality with the Language Localization pane. For more extensive language translation, go to Configuration > Remote Access VPN > Language Localization.

To change messages that appear on the AnyConnect GUI, perform the following steps:

Step 1 Click **Template** to expand the template area. Click **Export** to export the English language template to your local PC or a remote device.

Step 2 Edit the template and make changes to any messages. Edit any messages you want to change and save the file.The example below shows a message containing connection termination information:

```
msgid ""
"The VPN connection has been disconnected due to the system suspending. The
"reconnect capability is disabled.  A new connection requires re-"
"authentication and must be started manually.  Close all sensitive networked
"applications."
```

Step 3 Click **Import** to import the file you edited as a new translation template.

Step 4 Specify a language for the template. The template becomes a translation table in cache memory with the name you specify. Use an abbreviation that is compatible with the language options for your browser. For example, if you are creating a table for the Chinese language, and you are using IE, use the abbreviation *zh*, that is recognized by IE.

Step 5 Click **Apply** to make your changes to the security appliance.

Fields

Add—Launches the Add Localization Entry dialog where you can select a localization template to add and you can edit the contents of the template.

Edit—Launches the Edit Localization Entry dialog for the selected language in the table, and allows you to edit the previously-imported language localization table.

Delete—Deletes a selected language localization table.

Import—Launches the Import Language Localization dialog where you can import a language localization template or table.

Export—Launches the Export Language Localization dialog where you can export a language localization template or table to a URL where you can make changes to the table or template.

Language—The language of existing Language Localization tables.

Template—Expands the Template area:

- View—Displays the contents of the English language template.
- Export—Launches the Export Language Localization dialog where you can export the English language template to a URL where you can make changes.

Modes

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

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

Customized Installer Transforms

Specify transform files for customizing the AnyConnect client installation in this pane.

Note The security appliance does not support this feature for the AnyConnect VPN client, versions 2.0 and 2.1. For more information on manually customizing the client, *AnyConnect VPN Client Administrator Guide* and the *Release Notes for Cisco AnyConnect VPN Client*.

Fields

Import—Launches the Import AnyConnect Customization Objects dialog, where you can specify a transform file to import.

Export—Launches the Export AnyConnect Customization Objects dialog, where you can specify a transform file to export.

Delete—Removes the selected file.

Platform—The type of remote PC platform supported by the transform.

Object Name—The name of the transform.

Modes

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

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

Localized Installer Transforms

Specify files for localizing (translating) the AnyConnect client installer program in this pane.

Fields

Import—Launches the Import AnyConnect Customization Objects dialog, where you can specify a file to import as an transform.

Export—Launches the Export AnyConnect Customization Objects dialog, where you can specify a file to export as an transform.

Delete—Removes the selected transform.

Platform—The type of remote PC platform supported by the transform.

Object Name—The name of the transform.

Modes

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

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

Import/Export Language Localization

In the Import Translation Table and Export Translation Table dialog boxes you can import or export a translation table to the security appliance to provide translation of user messages.

Translation templates are XML files that contain message fields that can be edited with translated messages. You can export a template, edit the message fields, and import the template as a new translation table, or you can export an existing translation table, edit the message fields, and re-import the table to overwrite the previous version.

Fields

- Language—Enter a name for the language.

  When *exporting*, it is automatically filled-in with the name from the entry you selected in the table.

When *importing*, you enter the language name in the manner that you want it to be identified. The imported translation table then appears in the list with the abbreviation you designated. To ensure that your browser recognizes the language, use language abbreviations that are compatible with the language options of the browser. For example, if you are using IE, use ***zh*** as the abbreviation for the Chinese language.

- Localization Template Name—The name of the XML file containing the message fields. The following templates are available:
  - AnyConnect—Messages displayed on the user interface of the Cisco AnyConnect VPN Client.
  - CSD—Messages for the Cisco Secure Desktop (CSD).
  - customization—Messages on the logon and logout pages, portal page, and all the messages customizable by the user.
  - keepout—Message displayed to remote users when VPN access is denied.
  - PortForwarder—Messages displayed to Port Forwarding users.
  - url-list—Text that user specifies for URL bookmarks on the portal page.
  - webvpn—All the layer 7, AAA and portal messages that are not customizable.
  - plugin-ica—Messages for the Citrix plug-in.
  - plugin-rdp—Messages for the Remote Desktop Protocol plug-in.
  - plugin-telnet,ssh—Messages for the TELNET and SSH plug-in.
  - plugin-vnc—Messages for the VNC plug-in.
- Select a file—Choose the method by which you want to import or export the file.
  - Remote server—Select this option to import a customization file that resides on a remote server accessible from the security appliance.
  - Path—Identify the method to access the file (ftp, http, or https), and provide the path to the file.
  - Flash file system—Choose this method to export a file that resides on the security appliance.
  - Path—Provide the path to the file.
  - Browse Flash—Browse to the path for the file.
  - Local computer—Choose this method to import a file that resides on the local PC.
  - Path—Provide the path to the file.
  - Browse Local Files—Browse to the path for the file.
- Import/Export Now—Click to import or export the file.

Modes

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

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

Configuring Bookmarks

The Bookmarks panel lets you add, edit, delete, import, and export bookmark lists.

Use the Bookmarks panel to configure lists of servers and URLs for access over clientless SSL VPN. Following the configuration of a bookmark list, you can assign the list to one or more policies – group policies, dynamic access policies, or both. Each policy can have only one bookmark list. The list names populate a drop-down list on the URL Lists tab of each DAP.

Caution Configuring bookmarks does not prevent the user from visiting fraudulent sites or sites that violate your company's acceptable use policy. In addition to assigning a bookmark list to the group policy, dynamic access policy, or both, apply a web ACL to these policies to control access to traffic flows. Disable URL Entry on these policies to prevent user confusion over what is accessible. See Security Precautions, page 39-1 for instructions.

Fields

- Bookmarks—Displays the existing bookmark lists.
- Add—Click to add a new bookmark list.
- Edit—Click to edit the selected bookmark list.
- Delete—Click to delete the selected bookmark list.
- Import—Click to import a bookmark list.
- Export—Click to export a bookmark 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 Bookmark List

The Add/Edit Bookmark List dialog box configure lists of servers and URLs for access over lets you add, edit, or delete a URL list, and also order the items in a designated URL list.

Fields

- Bookmark List Name—Specifies the name of the list to be added or selects the name of the list to be modified or deleted.
- Bookmark Title—Specifies the URL name displayed to the user.
- URL—Specifies the actual URL associated with the display name.
- Add—Opens the Add Bookmark Entry dialog box, on which you can configure a new server or URL and display name.

- Edit—Opens the Edit Bookmark Entry dialog box, on which you can configure a new server or URL and display name.
- Delete—Removes the selected item from the URL list. There is no confirmation or undo.
- Move Up/Move Down—Changes the position of the selected item in the URL 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 Bookmark Entry

The Add Bookmark Entry dialog box lets you create a link or bookmark for a URL list.

Fields

- Bookmark Title—Enter a name for the bookmark to display for the user.
- URL (drop-down)—Use the drop-down menu to select the URL type: http, https, cifs, or ftp. The URL types of all imported plug-ins also populate this menu. Select the URL type of a plug-in if you want to display the plug-in as a link on the portal page.
- URL (text box)—Enter the DNS name or IP address for the bookmark. For a plug-in, enter the name of the server. Enter a forward slash and a question mark (/?) after the server name to specify optional parameters, then use an ampersand to separate parameter-value pairs, as shown in the following syntax:

  server/**?***Parameter*=*Value***&***Parameter*=*Value*

  For example:

  host/**?DesiredColor=4&DesiredHRes=1024&DesiredVRes=768**

  The particular plug-in determines the optional parameter-value pairs that you can enter.

  To provide single sign-on support for a plug-in, use the parameter-value pair **csco_sso=1**. For example:

  host/**?csco_sso=1&DesiredColor=4&DesiredHRes=1024&DesiredVRes=768**

  Note To access `\\server\share\subfolder\personal` *`folder`*, the user must have list permission for all points above `personal` *`folder`*.

- Subtitle—Provide additional user-visible text that describes the bookmark entry.
- Thumbnail—Use the drop-down menu to select an icon to associate with the bookmark on the end-user portal.
- Manage—Click to import or export images to use as thumbnails.

- Enable Smart Tunnel Option—Click to open the bookmark in a new window that uses the smart tunnel feature to pass data through the security appliance to or from the destination server. All browser traffic passes securely over the SSL VPN tunnel. This option lets you provide smart tunnel support for a browser-based application, whereas the Smart Tunnels option, also in the Clientless SSL VPN > Portal menu, lets you add nonbrowser-based applications to a smart tunnel list for assignment to group policies and usernames.
- Allow the users to bookmark the link—Check to let clientless SSL VPN users use the Bookmarks or Favorites options on their browsers. Uncheck to prevent access to these options.
- Advanced Options—(Optional) Open to configure further bookmark characteristics.
  - URL Method—Choose **Get** for simple data retrieval. Choose **Post** when processing the data might involve changes to it, for example, storing or updating data, ordering a product, or sending e-mail.
  - Post Parameters—Configure the particulars of the Post URL method.
  - Add/Edit—Click to add a post parameter.
  - Edit—Click to edit the highlighted post parameter.
  - Delete—Click to delete the highlighted post parameter.

Modes

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

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

Import/Export Bookmark List

You can import or export already configured bookmark lists. Import lists that are ready to use. Export lists to modify or edit them, and then reimport.

Fields

- Bookmark List Name—Identify the list by name. Maximum 64 characters, no spaces.
- Select a file—Choose the method by which you want to import or export the list file.
  - Local computer—Click to import a file that resides on the local PC.
  - Flash file system—Click to export a file that resides on the security appliance.
  - Remote server—Click to import a url list file that resides on a remote server accessible from the security appliance.
  - Path—Identify the method to access the file (ftp, http, or https), and provide the path to the file.
  - Browse Local Files/Browse Flash—Browse to the path for the file.
- Import/Export Now—Click to import or export the list file.

Modes

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

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

Configure GUI Customization Objects (Web Contents)

This dialogue box lets you import and export web content objects.

Fields

- File Name—Displays the names of the web content objects.
- File Type—Identifies the file type(s).
- Import/Export—Click to import or export a web content object.
- Delete—Click to delete the object.

Import/Export Web Content

Web contents can range from a wholly configured home page to icons or images you want to use when you customize the end user portal. You can import or export already configured web contents. Import web contents that are ready for use. Export web contents to modify or edit them, and then reimport.

Fields

- Source—Choose the location from which you want to import or export the file.
  - Local computer—Click to import or export a file that resides on the local PC.
  - Flash file system—Click to import or export a file that resides on the security appliance.
  - Remote server—Click to import a file that resides on a remote server accessible from the security appliance.
  - Path—Identify the method to access the file (ftp, http, or https), and provide the path to the file.
  - Browse Local Files.../Browse Flash...—Browse to the path for the file.
- Destination
  - Require authentication to access its content? Click **Yes** or **No**.
  - WebContent Path: Notice that the prefix to the path changes depending on whether you require authentication. The security appliance uses /+CSCOE+/ for objects that require authentication, and /+CSCOU+/ for objects that do not. The security appliance displays /+CSCOE+/ objects on the portal page only, while /+CSCOU+/ objects are visible and usable in either the logon or the portal pages.
- Import/Export Now—Click to import or export the 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/Edit Post Parameter

Use this pane to configure post parameters for bookmark entries and URL lists.

About Clientless SSL VPN Variable Substitutions

Clientless SSL VPN variables allow for substitutions in URLs and forms-based HTTP post operations. These variables, also known as macros, let you configure users for access to personalized resources that contain the user ID and password or other input parameters. Examples of such resources include bookmark entries, URL lists, and file shares.

Note For security reasons, password substitutions are disabled for file access URLs (cifs://).

Also for security reasons, use caution when introducing password substitutions for web links, especially for non-SSL instances.

Fields

- Name, Value—Provide the name and value of the parameters exactly as in the corresponding HTML form, for example: <input name="*param_name*" value="*param_value*">.

  You can choose one of the supplied variables from the drop-down list, or you can construct a variable. The variables you can choose from the drop-down list include the following:

No.	Variable Substitution	Definition
1	CSCO_WEBVPN_USERNAME	SSL VPN user login ID
2	CSCO_WEBVPN_PASSWORD	SSL VPN user login password
3	CSCO_WEBVPN_INTERNAL_PASSWORD	SSL VPN user internal resource password. This is a cached credential, and not authenticated by a AAA server. If a user enters this value, it is used as the password for auto-signon, instead of the password value.
4	CSCO_WEBVPN_CONNECTION_PROFILE	SSL VPN user login group drop-down, a group alias within the connection profile
5	CSCO_WEBVPN_MACRO1	Set via RADIUS/LDAP vendor-specific attribute. I f you are mapping this from LDAP via an ldap-attribute-map, the Cisco attribute that uses this variable is WEBVPN-Macro-Substitution-Value1. Variable substitution via RADIUS is performed by VSA#223.

No.	Variable Substitution	Definition
6	CSCO_WEBVPN_MACRO2	Set via RADIUS/LDAP vendor-specific attribute. If you are mapping this from LDAP via an ldap-attribute-map, the Cisco attribute that uses this variable is WEBVPN-Macro-Substitution-Value2. Variable substitution via RADIUS is performed by VSA#224.
7	CSCO_WEBVPN_PRIMARY_USERNAME	Primary user login ID for double authentication.
8	CSCO_WEBVPN_PRIMARY_PASSWORD	Primary user login password for double authentication.
9	CSCO_WEBVPN_SECONDARY_USERNAME	Secondary user login ID for double authentication.
10	CSCO_WEBVPN_SECONDARY_PASSWORD	Secondary user login ID for double authentication.

When the security appliance recognizes one of these six variable strings in an end-user request—in a bookmark or a post form—it replaces it with the user-specific value before passing the request to a remote server.

Note You can obtain the http-post parameters for any application by performing an HTTP Sniffer trace in the clear (without the security appliance involved). Here is a link to a free browser capture tool, also called an HTTP Analyzer: http://www.ieinspector.com/httpanalyzer/downloadV2/IEHttpAnalyzerV2.exe.

Using Variables 1 - 4

The security appliance obtains values for the first four substitutions from the SSL VPN Login page, which includes fields for username, password, internal password (optional), and group. It recognizes these strings in user requests, and replaces them with the value specific to the user before it passes the request on to a remote server.

For example, if a URL list contains the link, http://someserver/homepage/CSCO_WEBVPN_USERNAME.html, the security appliance translates it to the following unique links:

- For USER1 the link becomes http://someserver/homepage/USER1.html
- For USER2 the link is http://someserver/homepage/USER2.html

In the following case, cifs://server/users/CSCO_WEBVPN_USERNAME, lets the security appliance map a file drive to specific users:

- For USER1 the link becomes cifs://server/users/USER1
- For USER1 the link is cifs://server/users/USER2

Using Variables 5 and 6

Values for macros 5 and 6 are RADIUS or LDAP vendor-specific attributes (VSAs). These substitutions let you set substitutions configured on either a RADIUS or an LDAP server.

Using Variables 7 - 10

Each time the security appliance recognizes one of these four strings in an end-user request (a bookmark or a post form), it replaces it with the user-specific value before passing the request to a remote server.

Example 1: Setting a Homepage

The following example sets a URL for the homepage:

- WebVPN-Macro-Value1 (ID=223), type string, is returned as *wwwin-portal.example.com*
- WebVPN-Macro-Value2 (ID=224), type string, is returned as *401k.com*

To set a home page value, you would configure the variable substitution as

https://CSCO_WEBVPN_MACRO1, which would translate to https://wwwin-portal.example.com.

The best way to do this is to configure the Homepage URL parameter in ASDM.

Go to the Add/Edit Group Policy pane, from either the Network Client SSL VPN or Clientless SSL VPN Access section of ASDM. The paths are as follows:

- Configuration > Remote Access VPN > Network (Client) Access > Group Policies > Add/Edit Group Policy > Advanced > SSL VPN Client > Customization > Homepage URL attribute.
- Configuration > Remote Access VPN > Clientless SSL VPN Access > Group Policies > Add/Edit Group Policy > More Options > Customization > Homepage URL attribute.

Example 2: Setting a Bookmark or URL Entry

You can use an HTTP Post to log in to an OWA resource using an RSA one-time password (OTP) for SSL VPN authentication, and then the static, internal password for OWA e-mail access. The best way to do this is to add or edit a bookmark entry in ASDM.

There are several paths to the Add Bookmark Entry pane, including the following:

- Configuration > Remote Access VPN > Clientless SSL VPN Access > Portal > Bookmarks > Add/Edit Bookmark Lists > Add/Edit Bookmark Entry > Advanced Options area > Add/Edit Post Parameters (available after you click **Post** in the URL Method attribute).

  or

  (Available after you click **Post** in the URL Method attribute):
- Network (Client) Access > Dynamic Access Policies > Add/Edit Dynamic Access Policy > URL Lists tab > Manage button > Configured GUI Customization Objects > Add/Edit button > Add/Edit Bookmark List > Add/Edit Bookmark Entry > Advanced Options area > Add/Edit Post Parameters.

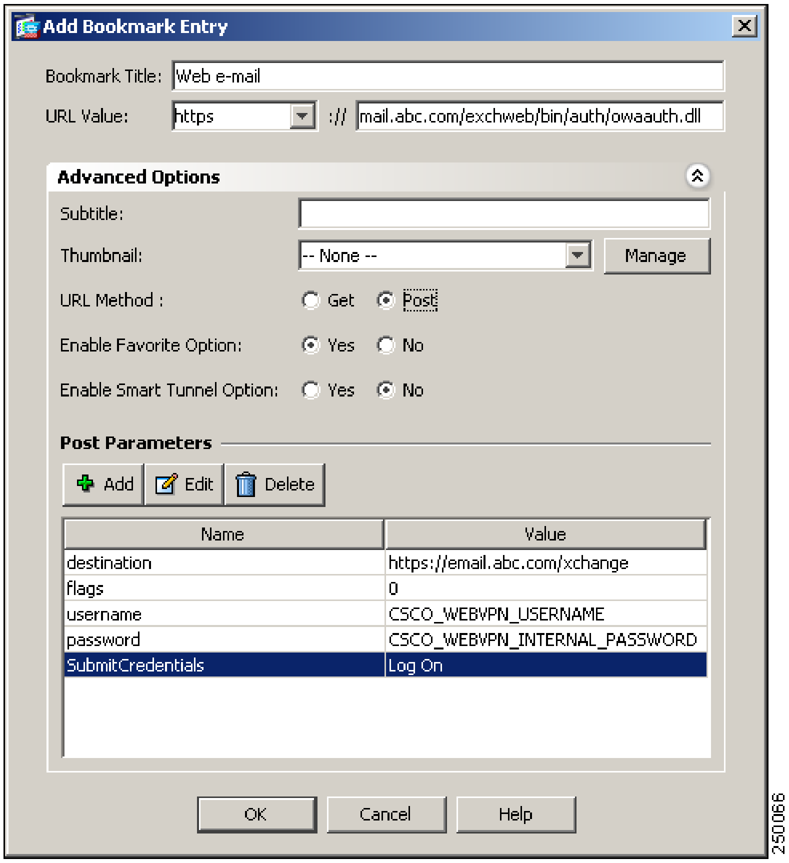

Add Bookmark Entry

Bookmark Title: Web e-mail

URL Value: https ://mail.abc.com/exchweb/bin/auth/owaauth.dll

Advanced Options

Subtitle:

Thumbnail: -- None --

URL Method : ○ Get ● Post

Enable Favorite Option: ● Yes ○ No

Enable Smart Tunnel Option: ○ Yes ● No

Post Parameters

Name	Value
destination	https://email.abc.com/xchange
flags	0
username	CSCO_WEBVPN_USERNAME
password	CSCO_WEBVPN_INTERNAL_PASSWORD
SubmitCredentials	Log On

Example 3: Configuring File Share (CIFS) URL Substitutions

You can allow a more flexible bookmark configuration by using variable substitution for CIFS URLs.

If you configure the URL cifs://server/CSCO_WEBVPN_USERNAME, the security appliance automatically maps it to the user’s file share home directory. This method also allows for password and internal password substitution. The following are example URL substitutions:

cifs://CSCO_WEBVPN_USERNAME:CSCO_WEBVPN_PASSWORD@server

cifs://CSCO_WEBVPN_USERNAME:CSCO_WEBVPN_INTERNAL_PASSWORD@server

cifs://domain;CSCO_WEBVPN_USERNAME:CSCO_WEBVPN_PASSWORD@server

cifs://domain;CSCO_WEBVPN_USERNAME:CSCO_WEBVPN_INTERNAL_PASSWORD@server

cifs://domain;CSCO_WEBVPN_USERNAME:CSCO_WEBVPN_PASSWORD@server/CSCO_WEBVPN_USERNAME

cifs://domain;CSCO_WEBVPN_USERNAME:CSCO_WEBVPN_INTERNAL_PASSWORD@server/CSCO_WEBVPN_USERNAME

More examples

For more variable substitution examples, see the *Cisco ASA 5500 SSL VPN Deployment Guide* on cisco.com.

Modes

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

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

CHAPTER 40

E-Mail Proxy

E-mail proxies extend remote e-mail capability to users of Clientless SSL VPN. 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 Clientless SSL VPN 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 Clientless SSL VPN 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 Clientless SSL VPN supports. By default, the Security Appliance listens to port 988, and connections automatically are 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

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

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

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.

DN Field	Definition
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

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.
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

The E-mail Proxy Access screen lets you identify interfaces on which to configure e-mail proxy. You can configure and edit e-mail proxies on individual interfaces, and you can configure and edit 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

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

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 Clientless SSL VPN 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 Clientless SSL VPN 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 Clientless SSL VPN 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

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

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 Clientless SSL VPN 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 41

Configuring SSL Settings

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 Clientless, browser-based 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 to negotiate 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 port forwarding for Clientless SSL VPN, 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 to negotiate as a client. 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
- **Certificates—**Lets you select a fallback certificate, and displays configured interfaces and the configured certificates associated with them.
  - **Fallback Certificate**—Click to select a certificate to use for interfaces that have no certificate associated with them. If you select **None**, the security appliance uses the default RSA key-pair and certificate.
  - **Interface** and **ID Certificate** columns—Display configured interfaces and the certificate, 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 Certificate

Fields

- **Interface**—Displays the name of the interface you are editing.

- **Certificate**—Click to select a previously enrolled certificate 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
•	•	•	•	—

SSL Certificates

In this pane, you can require that device management sessions require user certificates for SSL authentication.

Fields

- **Interface**—Displays the name of the interface you are editing.
- **User Certificate Required**—Click to select a previously enrolled certificate 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
•	•	•	•	—

PART 5

Monitoring the Security Appliance

CHAPTER 42

Monitoring Interfaces

ASDM lets you monitor interface statistics as well as interface-related features.

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 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

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

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

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

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

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 are 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

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

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, page 8-9, 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, page 8-9. 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

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

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 42-9
- Monitoring Statistics for, page 42-9

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

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 43

Monitoring VPN

This chapter describes how to use VPN monitoring parameters and statistics for the following:

- VPN statistics for specific Remote Access, LAN-to-LAN, Clientless SSL VPN, 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

Use this pane 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 pane when you click Show Graphs. This attribute is particularly useful when you want to clarify data in that pane before printing or exporting it. To change the title, choose 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 pane, choose the entry and click **Add**.
- Selected Graphs—Shows the types of tunnels selected.

  If you click Show Graphs, ASDM shows the active tunnels types listed in a single pane.

  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 column to the Selected Graphs column.

- Remove—Moves the selected tunnel type from the Selected Graphs column to the Available Graphs column.
- Show Graphs—Displays a pane consisting of graphs of the tunnel types displayed in the Selected Graphs column. Each type in the pane 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

Use this pane 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 panepane when you click Show Graphs. This attribute is particularly useful when you want to clarify data in that pane 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 pane, 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 pane.

  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 pane consisting of graphs of the session types displayed in the Selected Graphs box. Each type in the pane 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 panes show detailed parameters and statistics for a specific remote-access, LAN-to-LAN, Clientless SSL VPN, 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 choose. The detail tables show all the relevant parameters for each session.

Sessions

Use this pane to view session statistics for the adaptive security appliance.

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.
  - Site-to-Site—Shows the number of LAN-to-LAN sessions.
  - SSL VPN–Clientless—Shows the number of clientless browser-based VPN sessions.
- SSL VPN–With Client—Shows the number of client-based SSL VPN sessions. With ASA version 8.x and above , this represents the AnyConnect SSL VPN client 2.x and above.
  - SSL VPN–Inactive—Shows the number of SSL VPN sessions that are inactive on the remote computer.

Note An administrator can keep track of the number of users in the inactive state and can look at the statistics. The sessions that have been inactive for the longest time are marked as idle (and are automatically logged off) so that license capacity is not reached and new users can log in. You can also access these statististics using the **show vpn-sessiondb** CLI command (refer to the *Cisco Security Appliance Command Reference Guide)*.

  - SSL VPN–Total—Shows the number of client-based and clientless SSL VPN sessions.
  - E-mail Proxy—Shows the number of E-mail proxy sessions.
  - VPN Load Balancing—Shows the number of load-balanced VPN 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 IPSec 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, in this pane 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 (IPsec software and hardware clients) traffic.
  - Username/Connection Profile—Shows the username or login name and the connection profile (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.
  - Group Policy Connection Profile—Displays the tunnel group policy connection profile for the session.
  - 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.

Note The Assigned IP Address field does not apply to Clientless SSL VPN sessions, as the ASA (proxy) is the source of all traffic . For a hardward client session in Network Extension mode, the Assigned IP address is the subnet of the hardware client's private/inside network interface.

  - 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 (Peer) 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.
- IPSec Site-toSite—Indicates that the values in this table relate to LAN-to-LAN traffic.
  - Connection Profile/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.
- Clientless SSL VPN—Indicates that the values in this table relate to Clientless SSL VPN traffic.
  - Username/IP Address—Shows the username or login name for the session and the IP address of the client.
  - Group Policy Connection Profile—Displays the connection profile of the tunnel group policy.
  - 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.
- SSL VPN Client—Indicates that the values in this table relate to traffic for SSL VPN Client sessions.
  - Username/IP Address—Shows the username or login name for the session and the IP address of the client.
  - Group Policy Connection Profile—Displays the connection profile of the tunnel group policy.
  - 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.
- E-Mail Proxy—Indicates that the values in this table relate to traffic for Clientless SSL VPN sessions.
  - 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.
  - 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—Chooses a criterion to use to filter the sessions to be logged out. If you choose any but --All Sessions--, the box to the right of the Logout By list becomes active. If you choose the value Protocol for Logout By, the box becomes a list, from which you can choose 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 column.
- 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

The Session Details pane displays configuration settings, statistics, and state information about the selected session.

The Remote Detailed table at the top of the Session Details pane 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 panepane 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
•	—	•	—	—

Cluster Loads

Use this pane 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.
  - IPSec Load %—Indicates what percentage of a server's total capacity is in use, based upon the capacity of that server.
  - SSL Load %—Indicates what percentage of a SSL server's total capacity is in use, based upon the capacity of that server.
  - IPSec Sessions—Shows the number of currently active sessions.

  - SSL 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
•	—	•	—	—

Crypto Statistics

This pane 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.
- 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

This pane 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 clientless SSL VPN or SSL VPN Client sessions.
- 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
•	—	•	—	—

Encryption Statistics

This pane 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
•	—	•	—	—

Global IKE/IPSec Statistics

This pane 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.
- 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
•	—	•	—	—

NAC Session Summary

The NAC Session Summary pane 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

This pane 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 Tunnel—Shows the number of currently active sessions.
- Cumulative Tunnels—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
•	—	•	—	—

VLAN Mapping Sessions

This pane displays the number of sessions assigned to an egress VLAN, as determined by the value of the Restrict Access to VLAN parameter of each group policy in use. The security appliance forwards all traffic to the specified VLAN.

Field

- Active VLAN Mapping Sessions—Number of VPN sessions assigned to an egress VLAN.

Modes

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

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

SSO Statistics for Clientless SSL VPN Session

This pane displays the single sign-on statistics for currently active SSO servers configured for the security appliance.

Note These statistics are for SSO with SiteMinder and SAML Browser Post Profile servers only.

Fields

- Show Statistics For SSO Server—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 (SiteMinder servers)
  - Web Agent URL (SiteMinder servers)
  - Assertion Consumer URL (SAML POST servers)
  - Issuer (SAML POST servers)
  - 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
- Clear SSO Server Statistics—Resets statistics for the displayed server.

Modes

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

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

VPN Connection Status for the Easy VPN Client

Use this panel to view the status of the security appliance configured as an Easy VPN client. This features applies to the ASA5505 only.

Fields

VPN Client Detail—Displays configuration information for the ASA5505 configured as an Easy VPN Client.

Connect—Establishes a client connection

Refresh—Refreshes the information displayed in the VPN Client Detail panel.

Modes

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

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

CHAPTER 44

Monitoring Routing

You can use ASDM to monitor OSPF LSAs, OSPF and EIGRP neighbors, and the routing table. To access the routing monitoring screens, go to **Monitoring** > **Routing** in the ASDM interface.

This section contains the following topics:

- Monitoring OSPF LSAs, page 44-1
- Monitoring OSPF Neighbors, page 44-5
- Monitoring EIGRP Neighbors, page 44-7
- Displaying Routes, page 44-8

Monitoring OSPF LSAs

You can view the LSAs stored in the security appliance OSPF database in the **Monitoring** > **Routing** > **OSPF LSAs** area. 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

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

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

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

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

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

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
•	—	•	—	—

Monitoring OSPF Neighbors

The OSPF Neighbor pane displays the OSPF neighbors dynamically discovered and statically configured OSPF neighbors on the security appliance. The OSPF Neighbor pane is located at **Monitoring** > **Routing** >**OSPF Neighbors** in the ASDM interface.

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
•	—	•	—	—

Monitoring EIGRP Neighbors

The EIGRP Neighbors pane displays dynamically discovered EIGRP neighbors. Statically defined neighbors do not appear in this pane. To see the statically defined EIGRP neighbors, see **Configuration > Device Setup > Routing > EIGRP > Static Neighbor**.

Fields

- Address—IP address of the EIGRP neighbor.
- Interface—Interface on which the security appliance receives hello packets from the neighbor.
- Holdtime—Length of time (in seconds) that the security appliance waits to hear from the neighbor before declaring it down. This hold time is received from the neighbor in the hello packet, and begins decreasing until another hello packet is received from the neighbor.

  If the neighbor is using the default hold time, this number will be less than 15. If the peer configures a non-default hold time, the non-default hold time will be displayed.

  If this value reaches 0, the security appliance considers the neighbor unreachable.
- Uptime—Elapsed time (in hours:minutes: seconds) since the security appliance first heard from this neighbor.
- Queue Length—Number of EIGRP packets (update, query, and reply) that the security appliance is waiting to send.
- Sequence Number—Sequence number of the last update, query, or reply packet that was received from the neighbor.
- SRTT—Smooth round-trip time. This is the number of milliseconds required for an EIGRP packet to be sent to this neighbor and for the security appliance to receive an acknowledgment of that packet.
- RTO—Retransmission timeout (in milliseconds). This is the amount of time the security appliance waits before resending a packet from the retransmission queue to a neighbor.
- Clear Neighbors—Click the Clear Neighbors button to clear dynamically-learned neighbors from the 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
•	—	•	—	—

Displaying 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.
  - EIGRP—The route was derived using EIGRP.
  - 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 45

Monitoring Properties

This chapter includes the following sections:

Monitoring AAA Servers

This section includes the following topics:

Viewing AAA Server Statistics

Use this procedure to view statistics for AAA Servers.

Prerequisites

- You are connected to the security appliance using ASDM.
- You have already completed the initial security appliance configurations included in the ASDM Startup Wizard. For more information, see Using the Startup Wizard, page 7-1.
- You have already configured the servers and server groups that are being managed by the security appliance. For more information, see the Summary of Support, page 16-3.

Procedure

To view AAA Server statistics, perform the following steps.

Step 1 From the ASDM toolbar, click **Monitoring**.

The monitoring functions display in the left-hand Navigation pane.

Step 2 Click **Properties**.

The Properties Navigation pane opens.

Step 3 Click **AAA Servers**.

The AAA Servers dialog box opens in the right-hand pane, displaying a list of the configured AAA servers.

Step 4 Click the row for the server whose statistics you want to monitor.

Statistics for the selected server display in the lower portion of the dialog box.

See also:

- Fields Used to Monitor AAA Servers, page 45-3
- Updating the Operational State of an AAA Server, page 45-2
- Summary of Support, page 16-3
- AAA Overview, page 16-1

Updating the Operational State of an AAA Server

Use this procedure to update the operational state of an AAA server.

Prerequisites

- You are connected to the security appliance using ASDM.
- You have already completed the initial security appliance configurations included in the ASDM Startup Wizard. For more information, see Using the Startup Wizard, page 7-1.
- You have already configured the servers and server groups that are being managed by the security appliance. For more information, see the Summary of Support, page 16-3.

Procedure

To update the state of an AAA Server, perform the following steps.

Step 1 From the ASDM toolbar, click **Monitoring**.

The monitoring functions display in the left-hand Navigation pane.

Step 2 Click **Properties**.

The Properties Navigation pane opens.

Step 3 Click **AAA Servers**.

The AAA Servers dialog box opens in the right-hand pane, displaying a list of the AAA servers that are configured on the security appliance.

Step 4 Click the row for the server to update.

Statistics for the selected server display in the lower portion of the dialog box.

Step 5 Click **Update Server Statistics**.

The Update Server Statistics dialog box opens.

Step 6 From the AAA Server Status selection list, choose the operational state to apply to this server.

The security appliance is updated with the server current state.

Step 7 Click **OK**.

The dialog box closes.

See also:

- Fields Used to Monitor AAA Servers, page 45-3
- Viewing AAA Server Statistics, page 45-1
- Summary of Support, page 16-3
- AAA Overview, page 16-1

Fields Used to Monitor AAA Servers

The following table describes the fields for monitoring AAA Servers.

Field	Description
Server Group	The name of the server group where the server resides.
Protocol	The protocol used by the AAA server group.
IP Address	The IP address for the AAA server.
Status	The operational status of the AAA server. • Active • Failed

Field	Description
Statistics	The lower portion of the AAA Servers dialog box shows the following current information about the selected server: • Server port and/or hostname • Number of pending requests • Average round trip time • Number of authentication requests • Number of authorization requests • Number of accounting requests • Number of retransmissions • Number of accepts • Number of rejects • Number of challenges • Number of malformed responses • Number of bad authenticators • Number of timeouts • Number of unrecognized responses
Clear Server Statistics	Zeroes the counters for the selected server's statistics.
Update Server Status	Opens the Update Server Status dialog box for changing the operational state of the AAA server.
Refresh	Refreshes the dialog box display.

See also:

- Viewing AAA Server Statistics, page 45-1
- AAA Overview, page 16-1
- Summary of Support, page 16-3
- Viewing AAA Server Statistics, page 45-1

Modes

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

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

Monitoring Device Access

This section includes the following topics:

- Monitoring User Lockouts
- Monitoring Authenticated Users
- Monitoring Active Sessions
- Fields Used to Monitor Device Access

See also:

- Configuring Management Access Rules, page 18-15

Monitoring User Lockouts

This section includes the following topics:

- Viewing Lockouts, page 45-5
- Removing All User Lockouts, page 45-6
- Removing One User Lockout, page 45-7

See also:

- Fields for Monitoring User Lockouts, page 45-12
- Recovering from a Lockout, page 18-27

Viewing Lockouts

Use this procedure to view information about users who were locked out of the security appliance after failing to successfully authenticate with an AAA server.

Prerequisites

- You are connected to the security appliance using ASDM.
- You have already completed the initial security appliance configurations included in the ASDM startup wizard. For more information, see Using the Startup Wizard, page 7-1.
- You have already configured the servers and server groups that are being managed by the security appliance. For more information, see the Summary of Support, page 16-3.
- You have already configured the user accounts that are being managed by the security appliance Local server. For more information, see Adding a User Account, page 16-18.
- You have already configured authentication for the security appliance using the section, About Authentication, page 16-2.

Procedure

To view information about user lockouts, perform the following steps:

Step 1 From the ASDM toolbar, click **Monitoring**.

The monitoring functions display in the left-hand Navigation pane.

Step 2 Click **Properties**.

The Properties Navigation pane opens.

Step 3 Click the plus (+) symbol next to Device Access.

The list of Device Access functions expands below it.

Step 4 Click **AAA Local Locked Out Users**.

The AAA Local Locked Out Users dialog box opens in the right-hand pane, displaying a list of users who failed to successfully authenticate with an AAA server.

See also:

- Fields for Monitoring User Lockouts, page 45-12
- About Authentication, page 16-2
- Removing One User Lockout, page 45-7
- Removing All User Lockouts, page 45-6

Removing All User Lockouts

Use this procedure to remove the lockouts of all users who were locked out of the security appliance after failing to successfully authenticate with an AAA server.

Prerequisites

- You are connected to the security appliance using ASDM.
- You have already completed the initial security appliance configurations included in the ASDM startup wizard. For more information, see Using the Startup Wizard, page 7-1.
- You have already configured the servers and server groups that are being managed by the security appliance. For more information, see the Summary of Support, page 16-3.
- You have already configured the user accounts that are being managed by the security appliance Local server. For more information, see Adding a User Account, page 16-18.
- You have already configured authentication for the security appliance using the section, About Authentication, page 16-2.

Procedure

To clear all user lockouts from the security appliance, perform the following steps:

Step 1 From the ASDM toolbar, click **Monitoring**.

The monitoring functions display in the left-hand Navigation pane.

Step 2 Click **Properties**.

The Properties Navigation pane opens.

Step 3 Click the plus (+) symbol next to Device Access.

The list of Device Access functions expands below it.

Step 4 Click **AAA Local Locked Out Users**.

The AAA Local Locked Out Users dialog box opens in the right-hand pane, displaying a list of users who failed to successfully authenticate with an AAA server.

Step 5 Click **Refresh**.

The display is refreshed with current lockout information.

Step 6 Review the refreshed list to make sure that you want to remove all lockouts.

Step 7 Click **Clear All Lockouts**.

All lockouts from the security appliance are removed and usernames removed from the list.

See also:

- Fields for Monitoring User Lockouts, page 45-12
- About Authentication, page 16-2
- Removing One User Lockout, page 45-7

Removing One User Lockout

Use this procedure to remove a lockout for one user who was locked out of the security appliance after failing to successfully authenticate with an AAA server.

Prerequisites

- You are connected to the security appliance using ASDM.
- You have already completed the initial security appliance configurations included in the ASDM startup wizard. For more information, see Using the Startup Wizard, page 7-1.
- You have already configured the servers and server groups that are being managed by the security appliance. For more information, see the Summary of Support, page 16-3.
- You have already configured the user accounts that are being managed by the security appliance Local server. For more information, see Adding a User Account, page 16-18.
- You have already configured authentication for the security appliance using the section, About Authentication, page 16-2.

Procedure

To remove a user lockout, perform the following steps:

Step 1 From the ASDM toolbar, click **Monitoring**.

The monitoring functions display in the left-hand Navigation pane.

Step 2 Click **Properties**.

The Properties Navigation pane opens.

Step 3 Click the plus (+) symbol next to Device Access.

The list of Device Access functions expands below it.

Step 4 Click **AAA Local Locked Out Users**.

The AAA Local Locked Out Users dialog box opens in the right-hand pane, displaying a list of users who failed to successfully authenticate with an AAA server.

Step 5 Select the username from the list.

The row is highlighted.

Step 6 Click **Clear Selected Lockout**.

The lockout is removed for this user and the row is removed from the list.

See also:

- Fields for Monitoring User Lockouts, page 45-12
- About Authentication, page 16-2
- Viewing AAA Server Statistics, page 45-1

Monitoring Authenticated Users

Use this procedure to monitor users who have successfully authenticated with an AAA server.

Prerequisites

- You are connected to the security appliance using ASDM.
- You have already completed the initial security appliance configurations included in the ASDM startup wizard. For more information, see Using the Startup Wizard, page 7-1.
- You have already configured the servers and server groups that are being managed by the security appliance. For more information, see the Summary of Support, page 16-3.
- You have already configured the user accounts that are being managed by the security appliance Local server. For more information, see Adding a User Account, page 16-18.
- You have already configured authentication for the security appliance using the section, About Authentication, page 16-2.

Procedure

To monitor information about users who have successfully authenticated, perform the following steps:

Step 1 From the ASDM toolbar, click **Monitoring**.

The monitoring functions display in the left-hand Navigation pane.

Step 2 Click **Properties**.

The Properties Navigation pane opens.

Step 3 Click the plus (+) symbol next to Device Access.

The list of Device Access functions expands below it.

Step 4 Click **Authenticated Users**.

The Authenticated Users dialog box opens in the right-hand pane, displaying a list of users who have successfully authenticated with an AAA server.

See also:

- Fields for Monitoring Users Who Have Authenticated with a Server, page 45-13
- About Authentication, page 16-2
- Viewing AAA Server Statistics, page 45-1

Monitoring Active Sessions

This section includes the following procedures:

- Viewing Active Sessions, page 45-9
- Disconnecting an Active Session, page 45-11

See also

- Viewing Active Sessions, page 45-9

Viewing Active Sessions

Use this procedure to view the sessions that are currently connected to the security appliance.

Prerequisites

- You are connected to the security appliance using ASDM.
- You have already completed the initial security appliance configurations included in the ASDM startup wizard. For more information, see Using the Startup Wizard, page 7-1.
- You have already configured the servers and server groups that are being managed by the security appliance. For more information, see the Summary of Support, page 16-3.
- You have already configured the user accounts that are being managed by the security appliance Local server. For more information, see Adding a User Account, page 16-18.
- You have already configured the security appliance access for the session traffic you want to monitor. See the procedures in one of the following sections:
  - Configuring Device Access for ASDM, Telnet, or SSH, page 18-1
  - Configuring CLI Parameters, page 18-2

Procedure

To monitor active sessions, perform the following steps:

Step 1 From the ASDM toolbar, click **Monitoring**.

The monitoring functions display in the left-hand Navigation pane.

Step 2 Click **Properties**.

The Properties Navigation pane opens.

Step 3 Click the plus (+) symbol next to Device Access.

The list of Device Access functions expands below it.

Step 4 Click **ASDM/HTTPS/Telnet/SSH Sessions**.

A dialog box opens in the right-hand pane, displaying the list of currently active connections.

The following table describes the fields for monitoring active ASDM/HTTPS/Telnet sessions.

Field	Description
Type	The type of connection (ASDM/HTTPS/Telnet).
Session ID	The name of a currently connected ASDM/HTTPS/Telnet session.

Field	Description
IP Address	The IP address of the host or network that is currently connected to the security appliance.
Disconnect	Disconnects the selected ASDM/HTTPS/Telnet session from the security appliance.
Refresh	Refreshes the dialog box display.

The following table describes the fields for monitoring active SSH sessions.

Field	Description
Client	The client type for the selected SSH session.
User	The user name for the selected SSH session.
State	The state of the selected SSH session.
Version	The version of SSH used to connect to the security appliance.
Encryption (In)	The inbound encryption method used for the selected session.
Encryption (Out)	The outbound encryption method used for the selected session.
HMAC (In)	The configured HMAC for the selected inbound SSH session.
HMAC (Out)	The configured HMAC for the selected outbound SSH session.
SID	The session ID of the selected session.
Disconnect	Disconnects an active SSH session connected to the security appliance.
Refresh	Refreshes the dialog box display.

See also:

- Viewing Active Sessions, page 45-9
- Disconnecting an Active Session, page 45-11
- Fields Used to Monitor Device Access, page 45-12

Modes

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

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

See also:

- Fields Used to Monitor Device Access, page 45-12
- Disconnecting an Active Session, page 45-11

Disconnecting an Active Session

Use this procedure to disconnect an active ASDM/HTTPS, SSH, or Telnet session that is currently connected to the security appliance.

Prerequisites

- You are connected to the security appliance using ASDM.
- You have already completed the initial security appliance configurations included in the ASDM startup wizard. For more information, see Using the Startup Wizard, page 7-1.
- You have already configured the servers and server groups that are being managed by the security appliance. For more information, see the Summary of Support, page 16-3.
- You have already configured the user accounts that are being managed by the security appliance Local server. For more information, see Adding a User Account, page 16-18.
- You have already configured the security appliance access. See the procedures in one of the following sections:
  - Configuring Device Access for ASDM, Telnet, or SSH, page 18-1
  - Configuring CLI Parameters, page 18-2

Procedure

To disconnect an active security appliance session, perform the following steps:

Step 1 From the ASDM toolbar, click **Monitoring**.

The monitoring functions display in the left-hand Navigation pane.

Step 2 Click **Properties**.

The Properties Navigation pane opens.

Step 3 Click the plus (+) symbol next to Device Access.

The list of Device Access functions expands below it.

Step 4 Click **ASDM/HTTPS/Telnet/SSH Sessions**.

A dialog box opens in the right-hand pane, displaying a table which lists the currently active connections.

Step 5 In the table, select the session you want to disconnect.

The row is highlighted.

Step 6 Click **Disconnect**.

The session is disconnected from the security appliance, and removed from the table.

See also:

- Viewing Active Sessions, page 45-9
- Fields Used to Monitor Device Access, page 45-12

Fields Used to Monitor Device Access

This section includes the following topics:

- Fields for Monitoring User Lockouts, page 45-12
- Fields for Monitoring Users Who Have Authenticated with a Server, page 45-13

See also:

- Monitoring User Lockouts, page 45-5
- Monitoring Authenticated Users, page 45-8
- Monitoring Active Sessions, page 45-9

Fields for Monitoring User Lockouts

The following table describes the fields for monitoring locked out users.

Field	Description
Lock Time	The amount of time that the user has been locked out of the system.
Failed Attempts	The number of authentication attempts that the user failed.
User	A list of usernames of those users who are currently locked out of the security appliance because they were unable to successfully authenticate with the authentication server.
Clear Selected Lockout	Removes the lockout for the selected username and removes the username from the list.
Clear All Lockouts	Removes the lockout for all usernames in the list. **Note** We recommend that you refresh the list of locked out users and review it before clearing all lockouts.
Refresh	Refreshes the dialog box display.

See also:

- Monitoring User Lockouts, page 45-5
- About Authentication, page 16-2

Modes

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

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

Fields for Monitoring Users Who Have Authenticated with a Server

The following table describes the fields for monitoring authenticated users.

Field	Description
User	The usernames of users who have successfully authenticated with an authentication server.
IP Address	The IP addresses of users who have successfully authenticated with an authentication server.
Dynamic ACL	The dynamic access list of the user authenticated to use the security appliance.
Inactivity Timeout	The amount of time that the user connection must remain inactive before the session times out and the user is disconnected.
Absolute Timeout	The amount of time that the user can remain connected before the session closes and the user is disconnected.
Refresh	Refreshes the dialog box display.

See also:

- Monitoring Authenticated Users, page 45-8
- About Authentication, page 16-2

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 lets 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. This section includes the following topics:

- Perfmon
- Xlates

Perfmon

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

Fields

- Available Graphs—Lists the components you can graph.
  - AAA Perfmon—Displays the 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 Title—Shows the graph window name to which you want to add a graph type. To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title.
- Add—Click to move the selected entries in the Available Graphs list to the Selected Graphs list.
- Remove—Click to remove the selected statistic type from the Selected Graphs list.
- Show Graphs—Click to display a new or updated graph window.

Modes

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

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

Xlates

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

Fields

- Available Graphs—Lists the components you can graph.
  - Xlate Utilization—Displays the security appliance NAT utilization.
- Graph Window Title—Shows the graph window name to which you want to add a graph type. To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title.
- Add—Click to move the selected entries in the Available Graphs list to the Selected Graphs list.
- Remove—Click to remove the selected entry from the Selected Graphs list.
- Show Graphs—Click to display a new or updated graph window.

Modes

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

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

CRL

This pane allows you to view or clear associated CRLs of selected CA certificates.

Fields

- CA Certificate Name—Choose the name of the selected certificate from the drop-down list.
- View CRL—Click to view the selected CRL.
- Clear CRL—Click to 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
•	•	•	•	—

DNS Cache

The security appliance provides a local cache of DNS information from external DNS queries that are sent for certain clientless SSL VPN 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 because the size was exceeded or no more memory is available, the cache will be thinned by one third, based on the entries age. The oldest entries will be removed.

Fields

- Host— Shows the DNS name of the host.
- IP Address—Shows the address that resolves to the hostname.
- Permanent—Indicates whether the entry was made though a **name** command.
- Idle Time—Specifies the time elapsed since the security appliance last referred to that entry.
- Active—Indicates whether the entry has aged out. If there is not adequate space in cache, this entry may be deleted.
- Clear Cache—Click to clear the entire 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
•	•	•	•	—

IP Audit

The IP Audit pane lets you view the number of packets that match informational and attack signatures that are shown in graphical or tabular form. Each graph type shows the combined packets for all interfaces that have this feature enabled.

Fields

- Available Graphs—Lists the types of signatures available for monitoring. See IP Audit Signatures for 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—Click to add the selected graph type to the Selected Graphs list.
- Remove—Click to remove the selected graph type from the Selected Graphs list.
- Show Graphs—Click to display a new or updated graph window.
- Selected Graphs—Lists the graph types you want to show in the Selected Graphs list.

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

This pane lets you view the status of the security appliance memory, CPU, and block utilization. This section includes the following topics:

- Blocks
- CPU
- Memory

Blocks

This pane lets you view the free and used memory blocks. You can choose up to four types of statistics to show in one graph window. You can open multiple graph windows at the same time.

Fields

- Available Graphs —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 Title—Shows the graph window name to which you want to add a graph type. To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title.
- Add—Click to move the selected entries in the Available Graphs list to the Selected Graphs list.
- Remove—Click to remove the selected statistic type from the Selected Graphs list.
- Show Graphs—Click to display a new or updated graph window.

Modes

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

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

CPU

This pane lets you view the CPU utilization. You can choose up to four types of statistics to show in one graph window. You can open multiple graph windows at the same time.

Fields

- Available Graphs—Lists the components you can graph.
  - CPU Utilization—Displays the security appliance CPU utilization.
- Graph Window Title—Shows the graph window name to which you want to add a graph type. To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title.
- Add—Click to move the selected entries in the Available Graphs list to the Selected Graphs list.
- Remove—Click to remove the selected graph type from the Selected Graphs list.
- Show Graphs—Click to display a new or updated graph window.

Modes

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

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

Memory

This pane lets you view the memory utilization. You can choose up to four types of statistics to show in one graph window. You can open multiple graph windows at the same time.

Fields

- Available Graphs—Lists the components you can graph.
  - Free Memory—Displays the security appliance free memory.
  - Used Memory—Displays the security appliance used memory.
- Graph Window Title—Shows the graph window name to which you want to add a graph type. To use an existing window title, select one from the drop-down list. To display graphs in a new window, enter a new window title.
- Add—Click to move the selected entries in the Available Graphs list to the Selected Graphs list.
- Remove—Click to remove the selected graph type from the Selected Graphs list.
- Show Graphs—Click to display a new or updated graph window.

Modes

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

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

WCCP

The Web Cache Communication Protocol redirects IPv4 traffic flows to web caches in real-time. In ASDM, you can monitor packet redirection of an interface using WCCP. WCCP also provides load balancing, scaling, fault tolerance, and fail safe services. Load balancing is provided by hashing based on the destination IP address. The hash values are used to choose the egress interface for any traffic flow.

This protocol also enables the security appliance and WCCP clients to form service groups to support a service. This section includes the following topics:

- Service Groups
- Redirection

Service Groups

This pane allows you to view and refresh the service group, the display mode, and hash settings, which include the source and destination IP addresses and the source and destination port numbers.

Fields

- Service Group—Choose the applicable service group from the drop-down list.
- Display Mode—Choose the display mode from the drop-down list.
- Destination IP Address—Specify the destination IP address.
- Source IP Address—Specify the source IP address.
- Destination Port—Specify the destination port number.
- Source Port—Specify the source port number.
- WCCP Service Groups—*Display-only.* Shows the selected WCCP service group information.

  For example:

```
Global WCCP information:
    Router information:
    Router Identifier:                   -not yet determined-
    Protocol Version:                    2.0

    Service Identifier: web-cache
    Number of Cache Engines:             0
    Number of routers:                   0
    Total Packets Redirected:            0
    Redirect access-list:                -none-
    Total Connections Denied Redirect:   0
    Total Packets Unassigned:            0
    Group access-list:                   -none-
    Total Messages Denied to Group:      0
    Total Authentication failures:       0
    Total Bypassed Packets Received:     0

    Service Identifier: 1
    Number of Cache Engines:             0
    Number of routers:                   0
    Total Packets Redirected:            0
    Redirect access-list:                -none-
    Total Connections Denied Redirect:   0
    Total Packets Unassigned:            0
    Group access-list:                   -none-
    Total Messages Denied to Group:      0
    Total Authentication failures:       0
    Total Bypassed Packets Received:     0
```

Redirection

This pane allows you to view and refresh WCCP interface statistics in either a summary or detailed format.

Fields

- Show Summary—Choose this option to display statistics in a summary format.
- Show Details—Choose this option to display statistics in a detailed format.

- WCCP Interface Statistics—*Display-only*. Shows the current WCCP interface statistics.

  For example:

```
WCCP interface configuration details:
    Management0/0
    Output services: 0
    Input services:  1
    Static:          None
    Dynamic:         001
    Mcast services:  0
    Exclude In:      FALSE
```

CHAPTER 46

Monitoring Logging

You can view real-time syslog messages that appear in the log buffer. When you open the Cisco ASDM 6.2(1) for ASA 8.2(1) main application window, the most recent ASDM syslog messages appear at the bottom of a scrolling window. ASDM supports viewing of IPv6 addresses in syslog messages.

You can use these messages to help troubleshoot errors or monitor system usage and performance. For a description of the Logging feature, see Chapter 19, "Configuring Logging."

This chapter describes syslog message viewing, and includes the following sections:

- Log Buffer, page 46-1
- Real-Time Log Viewer, page 46-3

Log Buffer

The Log Buffer pane lets you view syslog messages that have been saved in the buffer in a separate window. To access this pane, choose **Monitoring** > **Logging** > **Log Buffer**.

To view the log buffer, perform the following steps:

Step 1 Choose the level of logging messages to view, ranging from Emergency to Debugging, from the drop-down list. For more information about severity levels, see Chapter 19, "Configuring Logging."

Step 2 Click **View** to open a separate pane in which log messages appear. To continue, see the "Log Buffer Viewer" section on page 46-2.

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

The Log Buffer Viewer pane lets you view messages that appear in the log buffer, an explanation of the message, details about the message, and recommended actions to take, if necessary, to resolve an error. To access this pane, choose **Monitoring** > **Logging** > **Log Buffer** > **View**.

To use the log buffer viewer, perform the following steps:

Step 1 Right-click a message to display a menu from which you can select from the Refresh, Copy Selected Log Entry, Save Log, Clear Display, Color Settings, Create Access Rule, Show Access Rule, and Show Details options. A list of icons associated with each severity level appears at the bottom of this pane.

Step 2 Choose from the following actions:

- Click **Refresh** to refresh the display.
- Click **Copy Selected Log Entry** to copy a selected message.
- Click **Save Log** to save the contents of the log to your computer.
- Click **Clear Display** to clear the list of messages.
- Click **Color Settings** to specify that messages of different severity levels display in different colors.
- Click **Create Access Rule** to create an access control rule that performs the opposite action of the access control rule that originally generated the message.
- Click **Show Access Rule** to show the access control rule that caused the selected message to be generated. This feature applies only to system log message IDs 106100 and 106023.
- Click **Show Details** to show or hide 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.
- In the Find field, enter text that you want to find in messages, and click **Search** to start the search.
- Click **Help** to obtain more information.
- Enter text to filter messages by in the Filter By drop-down list, then press **Enter** or click **Filter** to apply the filter to the displayed messages. Click **Show All** to display all messages. Filters are removed from the display. The Show All button is only active if a filter has been applied to the displayed syslog messages.

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

The Real-Time Log Viewer lets you view real-time syslog messages in a separate window. To access this pane, choose **Monitoring > Logging > Real-Time Log Viewer**.

To view syslog messages in real-time, perform the following steps:

Step 1 Choose the level of logging messages to view, ranging from Emergency to Debugging, from the drop-down list.

Step 2 Enter the buffer limit, which is the maximum number of syslog messages to view. The default is 1000.

Step 3 Click **View** to open a separate window in which syslog messages appear. To continue, see the "Real-Time Log Viewer" section on page 46-3.

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

The Real-Time Log Viewer pane lets you view incoming messages in real-time and filter them based on text that you specify. To access this pane, choose **Monitoring > Logging > Real-Time Log Viewer > View**.

To use the real-time log viewer, perform the following steps:

Step 1 Right-click a message in the viewer to display a menu from which you can select from the Pause, Copy Selected Log Entry, Save Log, Clear Display, Color Settings, Create Access Rule, Show Access Rule, and Show Details options. A list of color-coded icons that are associated with each severity level appears at the bottom of this pane. For more information about severity levels, see Chapter 19, "Configuring Logging."

Step 2 Choose from the following actions:

- Click **Pause** to stop the scrolling of the display.
- Click **Copy Selected Log Entry** to copy a selected message.
- Click **Save Log** to save the contents of the log to your computer.
- Click **Clear Display** to clear the list of messages.
- Click **Color Settings** to specify that messages of different severity levels display in different colors.
- Click **Create Access Rule** to create an access control rule that performs the opposite action of the access control rule that originally generated the message.

- Click **Show Access Rule** to show the access control rule that caused the selected message to be generated. This feature applies only to syslog message IDs 106100 and 106023.
- Click **Show Details** to show or hide 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.
- In the Find field, enter text that you want to find in messages, and click **Search** to start the search.
- Click **Help** to obtain more information.
- Enter text to filter messages by in the Filter By drop-down list, then press **Enter** or click **Filter** to apply the filter to the displayed messages. Click **Show All** to display all messages. Filters are removed from the display. The Show All button is only active if a filter has been applied to the displayed log 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 47

Monitoring Failover

Failover monitoring in ASDM depends upon the mode of the device. In single context mode, or within a security context in multiple context mode, you can monitor the state of failover for the device and view stateful failover statistics. In the system execution space of multiple context mode, you can monitor the failover state for each failover group

For more information about monitoring failover in each of these system configurations, see the following topics:

- Monitoring Failover in Single Context Mode or in a Security Context, page 47-1
- Monitoring Failover in the System Execution Space, page 47-6

Monitoring Failover in Single Context Mode or in a Security Context

You can monitor the status of the active and standby devices in a failover pair and failover related statistics in the **Monitoring** > **Properties** > **Failover** area. See the following screens for more information:

- Status—Displays the failover status of the device.
- Graphs—Displays graphs of various failover communication statistics.

Additional Information

For more information about failover in general, see Understanding Failover.

Status

The Status pane displays the failover state of the system. In single context mode you can control the failover state of the system by doing the following:

- Toggling the active/standby state of the device.
- Resetting a failed device.
- Reloading the standby unit.

In multiple context mode, you can control these settings in the system execution space. See Monitoring Failover in the System Execution Space, page 47-6.

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.

  When failover is disabled from the active peer after having been enabled, the standby peer shows Failover Off (pseudo-Standby), indicating that the standby peer continues to use its standby IP addresses even though it is no longer connected to an active peer. The standby peer continues to listen for a connection on its failover LAN. If failover is re-enabled on the active peer with a failover LAN configuration, then the standby peer resumes ordinary standby status after re-synchronizing the rest of its configuration. Otherwise, a pseudo-standby peer retains its status until it reloads or receives a command to become active or to re-enable failover.
- 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.
  - IPv6 ND tbl—The IPv6 Neighbor Discovery table 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—(Available in single mode and in system mode with multiple contexts) Click this button to make the security appliance the active unit in an active/standby configuration.
- Make Standby—(Available in single mode and in system mode with multiple contexts) Click this button to make the security appliance the standby unit in an active/standby pair.
- Reset Failover—(Available in single mode and in system mode with multiple contexts) 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—(Available in single mode and in system mode with multiple contexts) 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
•	•	•	—	•

Additional Information

For more information about failover in general, see Understanding 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
•	•	•	•	—

Additional Information

For more information about failover in general, see Understanding Failover.

Monitoring Failover in the System Execution Space

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

Additional Information

For more information about failover in general, see Understanding 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.
  - IPv6 ND tbl—the IPv6 Neighbor Discovery table 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
•	•	•	—	•

Additional Information

For more information about failover in general, see Understanding 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.
  - IPv6 ND tbl—the IPv6 Neighbor Discovery table 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 perform the 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
•	•	—	—	•

Additional Information

For more information about failover in general, see Understanding Failover.

CHAPTER 48

Monitoring Trend Micro Content Security

Note The ASA 5580 does not support the CSC SSM feature.

ASDM lets you monitor the CSC SSM statistics as well as CSC SSM-related features.

For an introduction to the CSC SSM, see the "Information About the CSC SSM" section on page 31-1.

Note If you have not completed the CSC 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 CSC Setup Wizard directly from Monitoring > Trend Micro Content Security.

Threats

To view information about various types of threats detected by the CSC SSM in a graph, perform the following steps:

Step 1 Choose **Monitoring > Trend Micro Content Security > Threats**.

The Available Graphs area lists the components whose statistics you can view in a graph. You can include a maximum of four graphs in one frame. The graphs display real-time data in 12-second intervals for the following:

- Viruses detected
- URLs filtered, URLs blocked
- Spam detected
- Files blocked
- Spyware blocked
- Damage Cleanup Services

Step 2 The Graph Window Title 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. The statistics already included in the graph window appear in the Selected Graphs list.

Step 3 To move the selected statistics type in the Available Graphs For list to the Selected Graphs list, click **Add**.

Step 4 To remove the selected statistics type from the Selected Graphs list, click **Remove**. The button name changes to **Delete** if the item you are removing was added from another pane, and is not being returned to the Available Graphs pane.

Step 5 To display a new window that shows a Graph tab and an updated graph with the selected statistics, click **Show Graphs**. Click the **Table** tab to display the same information in tabular form.

Step 6 From the Graph or Table tab, click **Export** in the menu bar or choose **File** > **Export** to save the graph or tabular information as a file on your local PC.

Step 7 From the Graph or Table tab, click **Print** in the menu bar or choose **File** > **Print** to print the information displayed in the window.

For more information, see the "Prerequisites for the CSC SSM" section on page 31-2.

Modes

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

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

Live Security Events

To view live, real-time security events in a separate window, perform the following steps:

Step 1 Choose **Monitoring** > **Trend Micro Content Security** > **Live Security Events**.

The Buffer Limit field shows the maximum number of log messages that you may view. The default is 1000.

Step 2 Click **View** to display the Live Security Events Log dialog box. You can pause incoming messages, clear the message window, and save event messages. You can also search messages for specific text.

For more information, see the "Information About the CSC SSM" section on page 31-1.

Modes

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

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

Live Security Events Log

To view live security events messages that are received from the CSC SSM, perform the following steps:

Step 1 To filter security event messages from the Filter By drop-down list, choose one of the following:

- Filter by Text, type the text, then click **Filter**.
- Show All, to display all messages or remove the filter.

Step 2 To use the Latest CSC Security Events pane, in which all columns are *display-only*, choose one of the following options:

- The time an event occurred.
- The IP address or hostname from which the threat came.
- The type of threat, or the security policy that determines event handling, or in the case of a URL filtering event, the filter that triggered the event.
- The subject of e-mails that include a threat, or the names of FTP files that include a threat, or blocked or filtered URLs.
- The recipient of e-mails that include a threat, or the IP address or hostname of a threatened node, or the IP address of a threatened client.
- The type of event (such as Web, Mail, or FTP), or the name of a user or group for HTTP or FTP events, which include a threat.
- The action taken upon the content of a message, such as cleaning attachments or deleting attachments.
- The action taken on a message, such as delivering it unchanged, delivering it after deleting the attachments, or delivering it after cleaning the attachments.

Step 3 To search security event messages based on the text that you enter, choose one of the following:

- In the Text field, enter the text to search for in the security event messages log, then click **Find Messages**.
- To find the next entry that matches the text you typed in this field, click **Find**.

Step 4 To pause the scrolling of the Live Security Events log, click **Pause**.

Step 5 To save the log to a file on your PC, click **Save**.

Step 6 To remove the list of messages, click **Clear Display**.

Step 7 To close the pane and return to the previous screen, click **Close**.

For more information, see the "Information About the CSC SSM" section on page 31-1.

Modes

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

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

Software Updates

To view information about CSC SSM software updates, perform the following steps:

Choose **Monitoring** > **Trend Micro Content Security** > **Software Updates**.

The Software Updates pane displays the following information, which is refreshed automatically about every 12 seconds:

- The names of parts of the CSC SSM software that can be updated.
- The current version of the corresponding component.
- The date and time that the corresponding component was last updated. If the component has not been updated since the CSC SSM software was installed, "None" appears in this column.
- The date and time that ASDM last received information about CSC SSM software updates.

For more information, see the "Information About the CSC SSM" section on page 31-1.

Modes

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

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

Resource Graphs

The adaptive security appliance lets you monitor CSC SSM status, including CPU resources and memory usage.

- CSC CPU, page 48-4
- CSC Memory, page 48-5

CSC CPU

To view CPU usage by the CSC SSM in a graph, perform the following steps:

Step 1 Choose **Monitoring** > **Trend Micro Content Security** > **Resource Graphs** > **CSC CPU**.

The CSC CPU pane displays the components whose statistics you can view in a graph, including statistics for CPU usage on the CSC SSM.

Step 2 To continue, go to Step 2 of the "Threats" section on page 48-1.

For more information, see the "Information About the CSC SSM" section on page 31-1.

Modes

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

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

CSC Memory

To view information about memory usage on the CSC SSM in a graph, perform the following steps:

Step 1 Choose **Monitoring** > **Trend Micro Content Security** > **Resource Graphs** > **CSC Memory**.

The Available Graphs area lists the components whose statistics you can view in a graph, including the following.

- The amount of memory not in use.
- The amount of memory in use.

Step 2 To continue, go to Step 2 of the "Threats" section on page 48-1.

For more information, see the "Information About the CSC SSM" section on page 31-1.

Modes

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

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

CHAPTER 49

Configuring Smart Call-Home

This chapter describes how to configure the Smart Call home feature. This chapter includes the following sections:

Information About Smart Call Home

Smart Call Home offers proactive diagnostics and real-time alerts on the security appliance and provides higher network availability and increased operational efficiency. Smart Call Home offers higher network availability through proactive, fast issue resolution by:

- Identifying issues quickly with continuous monitoring, real-time, proactive alerts, and detailed diagnostics.
- Making the customer aware of potential problems from Smart Call Home notification, a service request has been opened, with all diagnostic data attached.
- Resolving critical problems faster with direct, automatic access to experts at Cisco TAC.

Smart Call Home offers increased operational efficiency by providing customers the ability to:

- Use staff resources more efficiently by reducing troubleshooting time.
- Generate Service Requests to Cisco TAC automatically, routed to the appropriate support team, which provides detailed diagnostic information that speeds problem resolution.

Smart Call Home offers fast, web-based access to needed information that provides customers the ability to:

- Review all Call Home messages, diagnostics, and recommendations in one place.
- Check service request status quickly.
- View the most up-to-date inventory and configuration information for all Call Home devices.

Guidelines and Limitations

Failover Guidelines

Supports Active/Active and Active/Standby failover.

Firewall Mode Guidelines

Supported in routed and transparent firewall modes.

Context Mode Guidelines

Supported in single mode and multiple context mode.

IPv6 Guidelines

Support IPv6.

Additional Guidelines

In multiple context mode, the **snapshots** command is divided into two commands: one to obtain information from the system context and one to obtain information from the normal context.

Configuring Smart Call Home

You must configure smart call home for proactive diagnostics and real-time alerts for your security appliance.

Detailed Steps

Step 1 To configure Smart Call Home service, system setup, and alert subscription profiles, choose **Configuration> Device Management> Smart Call Home**.

Step 2 Check the **Enable Smart Call Home** check box to enable the feature.

Step 3 Double-click **System Setup**. The System Setup consists of three panes. Each pane can be expanded or collapsed by double-clicking the title row. In the Mail Servers pane. you can set up mail servers through which Smart Call Home messages are delivered to e-mail subscribers. In the Contact Information pane, you can enter the information of the person to contact for the security appliance that appears in Smart Call Home messages. This pane includes the following:

- The name of the contact person.
- The contact phone number.
- The postal address of the contact person.
- The e-mail address of the contact.
- The “from” e-mail address in Smart Call Home e-mail.
- The “reply-to” e-mail address in Smart Call Home e-mail.
- The customer ID.
- The site ID.
- The contract ID.

In the Alert Control pane, administrators can fine tune alert control parameters. This pane includes the following:

- The Alert group status pane, which lists the status (enabled or disabled) of alert groups.
- Enable and disable the diagnostics alert group.
- Enable and disable the configuration alert group.
- Enable and disable the environmental alert group.
- Enable and disable the inventory alert group.
- Enable and disable the snapshot alert group.
- Enable and disable the syslog alert group.
- Enable and disable the telemetry alert group.
- Enable and disable the threat alert group.
- The maximum number of Smart Call Home messages processed per minute.
- The “from” e-mail address in Smart Call Home e-mail.

Step 4 Double-click **Alert Subscription Profiles**. Each named subscription profile identifies subscribers and alert groups of interest.

a. Click **Add** or **Edit** to display the Subscription Profile Editor, in which you can create a new subscription profile or edit an existing subscription profile.

b. Click **Delete** to remove the selected profile.

c. Check the **Active** check box to send the Smart Call Home message of the selected subscription profile to subscribers.

Step 5 When you click **Add** or **Edit**, the Add or Edit Alert Subscription Profile dialog box appears.

a. The Name field is read-only, so you cannot edit it.

b. Check the **Enable this subscription profile** check box to enable or disable this particular profile.

c. Click either the **HTTP** or **Email** radio button in the Alert Delivery Method area.

d. In the Subscribers field, specify the alert delivery method: e-mail address or web address.

e. The Alert Dispatch area lets the administrator specify what type of Smart Call Home information to send to subscribers and under what conditions. There are two types of alerts, time-based and event-based, selected according to how the alert is triggered. The following alert groups are time-based: Configuration, Inventory, Snapshot, and Telemetry. The following alert groups are event-based: Diagnostic, Environmental, Syslog, and Threat.

f. The Message Parameters area lets you fine tune parameters that control messages sent to the subscriber, including the preferred message format and the maximum message size.

Step 6 For time-based alerts, in the Alert Dispatch area, click **Add** or **Edit** to display the Add or Edit Configuration Alert Dispatch Condition dialog box.

a. In the Alert Dispatch Frequency area, specify the frequency in which to send the information to subscribers:

- For monthly subscription, specify the day of the month, as well as the time of the day to send the information. If they are not specified, the security appliance chooses appropriate values for them.
- For weekly subscription, specify the day of the week, as well as the time of the day to send the information. If they are not specified, the security appliance chooses appropriate values for them.
- For daily subscription, specify the time of the day to send the information. If it is not specified, the security appliance chooses an appropriate value for it.

- For hourly subscription, specify the minute of the hour to send the information. If it is not specified, the security appliance chooses an appropriate value for it. Hourly subscription is only applicable to the snapshot and telemetry alert groups.

b. Click the **Basic** or **Detailed** radio button to provide the desired level of information to subscribers.

c. Click **OK** when you are done.

Step 7 For diagnostic, environment, and threat event-based alerts, in the Alert Dispatch area, click **Add** or **Edit** to display the Create or Edit Diagnostic Alert Dispatch Condition dialog box.

Step 8 Specify the event severity that triggers dispatch of the alert to subscribers in the Event Severity drop-down list, and then click **OK**.

Step 9 For inventory time-based alerts, in the Alert Dispatch area, click **Add** or **Edit** to display the Create or Edit Inventory Alert Dispatch Condition dialog box.

Step 10 Specify how often to dispatch alerts to subscribers in the Alert Dispatch Frequency drop-down list, and then click **OK**.

Step 11 For snapshot time-based alerts, in the Alert Dispatch area, click **Add** or **Edit** to display the Create or Edit Snapshot Alert Dispatch Condition dialog box.

a. In the Alert Dispatch Frequency area, specify the frequency in which to send the information to subscribers:

   – For monthly subscription, specify the day of the month, as well as the time of the day to send the information. If they are not specified, the security appliance chooses appropriate values for them.

   – For weekly subscription, specify the day of the week, as well as the time of the day to send the information. If they are not specified, the security appliance chooses appropriate values for them.

   – For daily subscription, specify the time of the day to send the information. If it is not specified, the security appliance chooses an appropriate value for it.

   – For hourly subscription, specify the minute of the hour to send the information. If it is not specified, the security appliance chooses an appropriate value for it. Hourly subscription is only applicable to the snapshot and telemetry alert groups.

   – For interval subscription, specify how often, in minutes, the formation is sent to the subscribers. This requirement is only applicable to the snapshot alert group.

b. Click **OK** when you are done.

Step 12 For syslog event-based alerts, in the Alert Dispatch area, click **Add** or **Edit** to display the Create or Edit Syslog Alert Dispatch Condition dialog box.

a. Check the **Specify the event severity which triggers the dispatch of alert to subscribers check box**, and choose the event severity from the drop-down list.

b. Check the **Specify the message IDs of syslogs which trigger the dispatch of alert to subscribers** check box.

c. Specify the syslog message IDs that trigger dispatch of the alert to subscribers according to the on-screen instructions.

d. Click **OK** when you are done.

Step 13 For telemetry event-based alerts, in the Alert Dispatch area, click **Add** or **Edit** to display the Create or Edit Telemetry Alert Dispatch Condition dialog box.

a. In the Alert Dispatch Frequency area, specify the frequency in which to send the information to subscribers:

- For monthly subscription, specify the day of the month, as well as the time of the day to send the information. If they are not specified, the security appliance chooses appropriate values for them.
- For weekly subscription, specify the day of the week, as well as the time of the day to send the information. If they are not specified, the security appliance chooses appropriate values for them.
- For daily subscription, specify the time of the day to send the information. If it is not specified, the security appliance chooses an appropriate value for it.
- For hourly subscription, specify the minute of the hour to send the information. If it is not specified, the security appliance chooses an appropriate value for it. Hourly subscription is only applicable to the snapshot and telemetry alert groups.

b. Click **OK** when you are done.

Feature History for Smart Call Home

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

Table 1 Feature History for Smart Call Home

Feature Name	Platform Releases	Feature Information
Smart Call Home	8.2(2)	The Smart Call Home feature offers proactive diagnostics and real-time alerts on the security appliance and provides higher network availability and increased operational efficiency. The following screens were introduced or modified: **Configuration> Device Management> Smart Call Home**

PART 6

Reference

APPENDIX A

Troubleshooting

This chapter describes how to troubleshoot the security appliance, and includes the following sections:

- Testing Your Configuration, page A-1
- Reloading the Security Appliance, page A-6
- Performing Password Recovery, page A-7
- Using the ROM Monitor to Load a Software Image, page A-10
- Erasing the Flash File System, page A-11
- Other Troubleshooting Tools, page A-12
- Common Problems, page A-13

Testing Your Configuration

This section describes how to test connectivity for the single mode security appliance or for each security context, how to ping the security appliance interfaces, and how to allow hosts on one interface to ping through to hosts on another interface.

We recommend that you only enable pinging and debug messages during troubleshooting. When you are done testing the security appliance, follow the steps in "Disabling the Test Configuration" section on page A-5.

This section includes the following topics:

- Enabling ICMP Debug Messages and System Log Messages, page A-1
- Pinging Security Appliance Interfaces, page A-2
- Pinging Through the Security Appliance, page A-4
- Disabling the Test Configuration, page A-5

Enabling ICMP Debug Messages and System Log Messages

Debug messages and system log messages can help you troubleshoot why your pings are not successful. The security appliance only shows ICMP debug messages for pings to the security appliance interfaces, and not for pings through the security appliance to other hosts. To enable debugging and system log messages, perform the following steps:

Step 1 To show ICMP packet information for pings to the security appliance interfaces, enter the following command:

```
hostname(config)# debug icmp trace
```

Step 2 To set system log messages to be sent to Telnet or SSH sessions, enter the following command:

```
hostname(config)# logging monitor debug
```

You can alternately use the **logging buffer debug** command to send log messages to a buffer, and then view them later using the **show logging** command.

Step 3 To send the system log messages to a Telnet or SSH session, enter the following command:

```
hostname(config)# terminal monitor
```

Step 4 To enable system log messages, enter the following command:

```
hostname(config)# logging on
```

The following example shows a successful ping from an external host (209.165.201.2) to the security appliance outside interface (209.165.201.1):

```
hostname(config)# debug icmp trace
Inbound ICMP echo reply (len 32 id 1 seq 256) 209.165.201.1 > 209.165.201.2
Outbound ICMP echo request (len 32 id 1 seq 512) 209.165.201.2 > 209.165.201.1
Inbound ICMP echo reply (len 32 id 1 seq 512) 209.165.201.1 > 209.165.201.2
Outbound ICMP echo request (len 32 id 1 seq 768) 209.165.201.2 > 209.165.201.1
Inbound ICMP echo reply (len 32 id 1 seq 768) 209.165.201.1 > 209.165.201.2
Outbound ICMP echo request (len 32 id 1 seq 1024) 209.165.201.2 > 209.165.201.1
Inbound ICMP echo reply (len 32 id 1 seq 1024) 209.165.201.1 > 209.165.201.2
```

This example shows the ICMP packet length (32 bytes), the ICMP packet identifier (1), and the ICMP sequence number (the ICMP sequence number starts at 0 and is incremented each time that a request is sent).

Pinging Security Appliance Interfaces

To test whether the security appliance interfaces are up and running and that the security appliance and connected routers are operating correctly, you can ping the security appliance interfaces. To ping the security appliance interfaces, perform the following steps:

Step 1 Draw a diagram of your single-mode security appliance or security context that shows the interface names, security levels, and IP addresses.

Note Although this procedure uses IP addresses, the **ping** command also supports DNS names and names that are assigned to a local IP address with the **name** command.

The diagram should also include any directly connected routers, and a host on the other side of the router from which you will ping the security appliance. You will use this information in this procedure and in the procedure in "Pinging Through the Security Appliance" section on page A-4. For example:

Figure A-1 Network Diagram with Interfaces, Routers, and Hosts

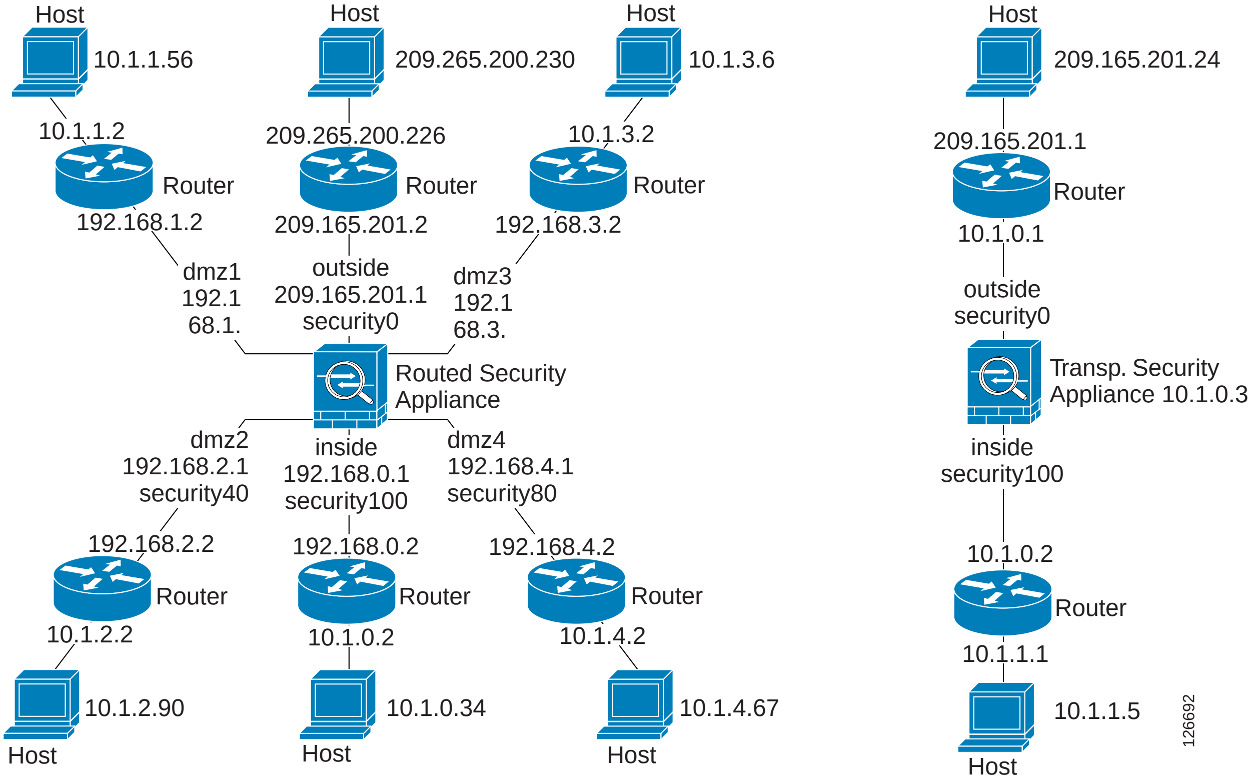

Step 2 Ping each security appliance interface from the directly connected routers. For transparent mode, ping the management IP address. This test ensures that the security appliance interfaces are active and that the interface configuration is correct.

A ping might fail if the security appliance interface is not active, the interface configuration is incorrect, or if a switch between the security appliance and a router is down (see Figure A-2). In this case, no debug messages or system log messages appear, because the packet never reaches the security appliance.

Figure A-2 Ping Failure at Security Appliance Interface

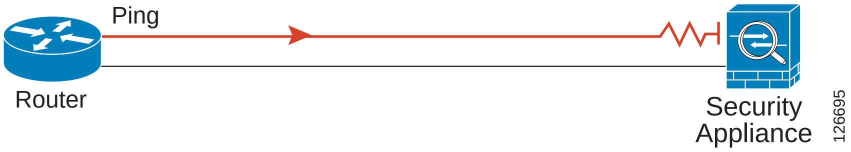

If the ping reaches the security appliance, and the security appliance responds, debug messages similar to the following appear:

```
ICMP echo reply (len 32 id 1 seq 256) 209.165.201.1 > 209.165.201.2
ICMP echo request (len 32 id 1 seq 512) 209.165.201.2 > 209.165.201.1
```

If the ping reply does not return to the router, then a switch loop or redundant IP addresses may exist (see Figure A-3).

Figure A-3 Ping Failure Because of IP Addressing Problems

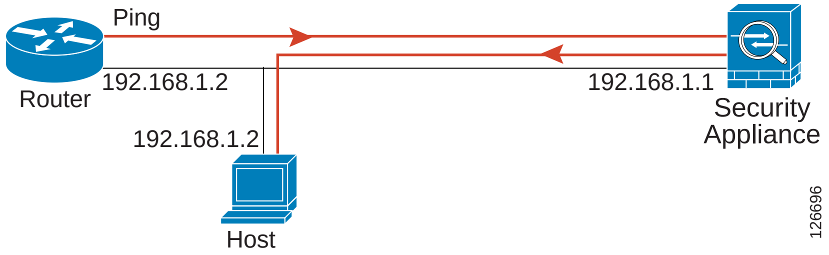

Step 3 Ping each security appliance interface from a remote host. For transparent mode, ping the management IP address. This test checks whether the directly connected router can route the packet between the host and the security appliance, and whether the security appliance can correctly route the packet back to the host.

A ping might fail if the security appliance does not have a return route to the host through the intermediate router (see Figure A-4). In this case, the debug messages show that the ping was successful, but system log message 110001 appears, indicating a routing failure.

Figure A-4 Ping Failure Because the Security Appliance has no Return Route

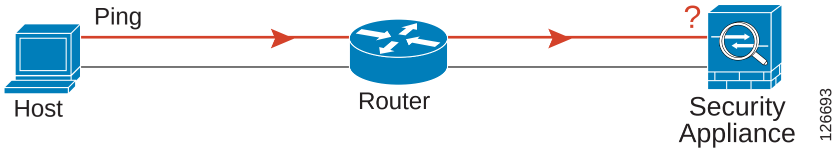

Pinging Through the Security Appliance

After you successfully ping the security appliance interfaces, make sure traffic can pass successfully through the security appliance. For routed mode, this test shows that NAT is operating correctly, if configured. For transparent mode, which does not use NAT, this test confirms that the security appliance is operating correctly. If the ping fails in transparent mode, contact Cisco TAC.

To ping between hosts on different interfaces, perform the following steps:

Step 1 To add an access list allowing ICMP from any source host, enter the following command:

```
hostname(config)# access-list ICMPACL extended permit icmp any any
```

By default, when hosts access a lower security interface, all traffic is allowed through. However, to access a higher security interface, you need the preceding access list.

Step 2 To assign the access list to each source interface, enter the following command:

```
hostname(config)# access-group ICMPACL in interface interface_name
```

Repeat this command for each source interface.

Step 3 To enable the ICMP inspection engine and ensure that ICMP responses may return to the source host, enter the following commands:

```
hostname(config)# class-map ICMP-CLASS
hostname(config-cmap)# match access-list ICMPACL
hostname(config-cmap)# policy-map ICMP-POLICY
hostname(config-pmap)# class ICMP-CLASS
hostname(config-pmap-c)# inspect icmp
hostname(config-pmap-c)# service-policy ICMP-POLICY global
```

Alternatively, you can also apply the ICMP access list to the destination interface to allow ICMP traffic back through the security appliance.

Step 4 Ping from the host or router through the source interface to another host or router on another interface.

Repeat this step for as many interface pairs as you want to check.

If the ping succeeds, a system log message appears to confirm the address translation for routed mode (305009 or 305011) and that an ICMP connection was established (302020). You can also enter either the **show xlate** or **show conns** command to view this information.

If the ping fails for transparent mode, contact Cisco TAC.

For routed mode, the ping might fail because NAT is not configured correctly (see Figure A-5). This failure is more likely to occur if you enable NAT control. In this case, a system log message appears, showing that the NAT failed (305005 or 305006). If the ping is from an outside host to an inside host, and you do not have a static translation (required with NAT control), the following system log message appears: “106010: deny inbound icmp.”

Note The security appliance only shows ICMP debug messages for pings to the security appliance interfaces, and not for pings through the security appliance to other hosts.

Figure A-5 Ping Failure Because the Security Appliance is not Translating Addresses

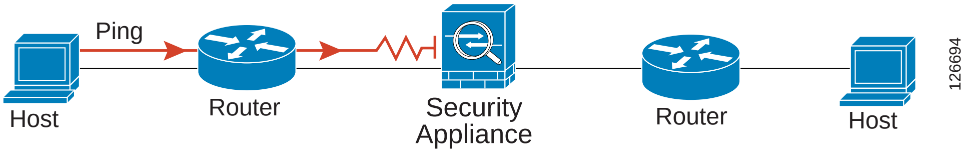

Disabling the Test Configuration

After you complete your testing, disable the test configuration that allows ICMP to and through the security appliance and that prints debug messages. If you leave this configuration in place, it can pose a serious security risk. Debug messages also slow the security appliance performance.

To disable the test configuration, perform the following steps:

Step 1 To disable ICMP debug messages, enter the following command:

```
hostname(config)# no debug icmp trace
```

Step 2 To disable logging, if desired, enter the following command:

```
hostname(config)# no logging on
```

Step 3 To remove the ICMPACL access list, and delete the related **access-group** commands, enter the following command:

```
hostname(config)# no access-list ICMPACL
```

Step 4 (Optional) To disable the ICMP inspection engine, enter the following command:

```
hostname(config)# no service-policy ICMP-POLICY
```

Traceroute

You can trace the route of a packet using the traceroute feature, which is accessed with the **traceroute** command. A traceroute works by sending UDP packets to a destination on an invalid port. Because the port is not valid, the routers along the way to the destination respond with an ICMP Time Exceeded Message, and report that error to the security appliance.

For more information, see Determining Packet Routing with Traceroute.

Packet Tracer

In addition, you can trace the lifespan of a packet through the security appliance to see whether the packet is operating correctly with the packet tracer tool. This tool lets you do the following:

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

The **packet-tracer** command provides detailed information about the packets and how they are processed by the security appliance. If a command from the configuration did not cause the packet to drop, the **packet-tracer** command will provide information about the cause in an easily readable manner. For example, when a packet is dropped because of an invalid header validation, the following message appears: “packet dropped due to bad ip header (reason).”

For more information, see Tracing Packets with Packet Tracer.

Reloading the Security Appliance

In multiple mode, you can only reload from the system execution space. To reload the security appliance, enter the following command:

```
hostname# reload
```

Recovering from a Lockout

In some circumstances, when you turn on command authorization or CLI authentication, you can be locked out of the security appliance CLI. You can usually recover access by restarting the security appliance. For information on common lockout conditions and how you might recover from them, see

Performing Password Recovery

This section describes how to recover passwords if you have forgotten them or you are locked out because of AAA settings, and how to disable password recovery for extra security. This section includes the following topics:

- Recovering Passwords for the ASA 5500 Series Adaptive Security Appliance, page A-7
- Recovering Passwords for the PIX 500 Series Security Appliance, page A-8
- Disabling Password Recovery, page A-9
- Using the ROM Monitor to Load a Software Image, page A-10

Recovering Passwords for the ASA 5500 Series Adaptive Security Appliance

To recover passwords for the ASA 5500 Series adaptive security appliance, perform the following steps:

Step 1 Connect to the adaptive security appliance console port.

Step 2 Power off the adaptive security appliance, and then power it on.

Step 3 After startup, press the **Escape** key when you are prompted to enter ROMMON mode.

Step 4 To update the configuration register value, enter the following command:

```
rommon #1> confreg 0x41
Update Config Register (0x41) in NVRAM...
```

Step 5 To set the adaptive security appliance to ignore the startup configuration, enter the following command:

```
rommon #1> confreg
```

The adaptive security appliance displays the current configuration register value, and asks whether you want to change it:

```
Current Configuration Register: 0x00000041
Configuration Summary:
  boot default image from Flash
  ignore system configuration

Do you wish to change this configuration? y/n [n]: y
```

Step 6 Record the current configuration register value, so you can restore it later.

Step 7 At the prompt, enter **Y** to change the value.

The adaptive security appliance prompts you for new values.

Step 8 Accept the default values for all settings. At the prompt, enter **Y**.

Step 9 Reload the adaptive security appliance by entering the following command:

```
rommon #2> boot
Launching BootLoader...
Boot configuration file contains 1 entry.

Loading disk0:/asa800-226-k8.bin... Booting...Loading...
```

The adaptive security appliance loads the default configuration instead of the startup configuration.

Step 10 Access the privileged EXEC mode by entering the following command:

```
hostname> enable
```

Step 11 When prompted for the password, press **Enter**.

The password is blank.

Step 12 Access the global configuration mode by entering the following command:

```
hostname# configure terminal
```

Step 13 Change the passwords, as required, in the default configuration by entering the following commands:

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

Step 14 Load the default configuration by entering the following command:

```
hostname(config)# no config-register
```

The default configuration register value is 0x1. For more information about the configuration register, see the *Cisco ASA 5500 Series Command Reference*.

Step 15 Save the new passwords to the startup configuration by entering the following command:

```
hostname(config)# copy running-config startup-config
```

Recovering Passwords for the PIX 500 Series Security Appliance

Recovering passwords on the PIX 500 series security appliance erases the login password, enable password, and **aaa authentication console** commands. To recover passwords for the PIX 500 series security appliance, perform the following steps:

Step 1 Download the PIX password tool from Cisco.com to a TFTP server accessible from the security appliance. For instructions, go to the following URL:

http://www.cisco.com/en/US/products/hw/vpndevc/ps2030/products_password_recovery09186a008009478b.shtml

Step 2 Connect to the security appliance console port.

Step 3 Power off the security appliance, and then power it on.

Step 4 Immediately after the startup messages appear, press the **Escape** key to enter monitor mode.

Step 5 In monitor mode, configure the interface network settings to access the TFTP server by entering the following commands:

```
monitor> interface interface_id
monitor> address interface_ip
monitor> server tftp_ip
monitor> file pw_tool_name
monitor> gateway gateway_ip
```

Step 6 Download the PIX password tool from the TFTP server by entering the following command:

```
monitor> tftp
```

If you have trouble reaching the server, enter the **ping** *address* command to test the connection.

Step 7 At the “Do you wish to erase the passwords?” prompt, enter **Y**.

You can log in with the default login password of “cisco” and the blank enable password.

The following example shows password recovery on a PIX 500 series security appliance with the TFTP server on the outside interface:

```
monitor> interface 0
0: i8255X @ PCI(bus:0 dev:13 irq:10)
1: i8255X @ PCI(bus:0 dev:14 irq:7 )

Using 0: i82559 @ PCI(bus:0 dev:13 irq:10), MAC: 0050.54ff.82b9
monitor> address 10.21.1.99
address 10.21.1.99
monitor> server 172.18.125.3
server 172.18.125.3
monitor> file np70.bin
file np52.bin
monitor> gateway 10.21.1.1
gateway 10.21.1.1
monitor> ping 172.18.125.3
Sending 5, 100-byte 0xf8d3 ICMP Echoes to 172.18.125.3, timeout is 4 seconds:
!!!!!
Success rate is 100 percent (5/5)
monitor> tftp
tftp np52.bin@172.18.125.3 via 10.21.1.1
Received 73728 bytes

Cisco PIX password tool (4.0) #0: Tue Aug 22 23:22:19 PDT 2005
Flash=i28F640J5 @ 0x300
BIOS Flash=AT29C257 @ 0xd8000

Do you wish to erase the passwords? [yn] y
Passwords have been erased.

Rebooting....
```

Disabling Password Recovery

You might want to disable password recovery to ensure that unauthorized users cannot use the password recovery mechanism to compromise the security appliance. To disable password recovery, enter the following command:

```
hostname(config)# no service password-recovery
```

On the ASA 5500 series adaptive security appliance, the **no service password-recovery** command prevents a user from entering ROMMON mode with the configuration intact. When a user enters ROMMON mode, the security appliance prompts the user to erase all Flash file systems. The user cannot enter ROMMON mode without first performing this erasure. If a user chooses not to erase the Flash file system, the security appliance reloads. Because password recovery depends on using ROMMON mode and maintaining the existing configuration, this erasure prevents you from recovering a password. However, disabling password recovery prevents unauthorized users from viewing the configuration or inserting different passwords. In this case, to restore the system to an operating state, load a new image and a backup configuration file, if available.

The **service password-recovery** command appears in the configuration file for information only. When you enter the command at the CLI prompt, the setting is saved in NVRAM. The only way to change the setting is to enter the command at the CLI prompt. Loading a new configuration with a different version of the command does not change the setting. If you disable password recovery when the security appliance is configured to ignore the startup configuration at startup (in preparation for password recovery), then the security appliance changes the setting to load the startup configuration as usual. If you use failover, and the standby unit is configured to ignore the startup configuration, then the same change is made to the configuration register when the **no service password recovery** command replicates to the standby unit.

On the PIX 500 series security appliance, the **no service password-recovery** command forces the PIX password tool to prompt the user to erase all Flash file systems. The user cannot use the PIX password tool without first performing this erasure. If a user chooses not to erase the Flash file system, the security appliance reloads. Because password recovery depends on maintaining the existing configuration, this erasure prevents you from recovering a password. However, disabling password recovery prevents unauthorized users from viewing the configuration or inserting different passwords. In this case, to restore the system to an operating state, load a new image and a backup configuration file, if available.

Using the ROM Monitor to Load a Software Image

This section describes how to load a software image to an adaptive security appliance from the ROM monitor mode using TFTP.

To load a software image to an adaptive security appliance, perform the following steps:

Step 1 Connect to the adaptive security appliance console port.

Step 2 Power off the adaptive security appliance, and then power it on.

Step 3 During startup, press the **Escape** key when you are prompted to enter ROMMON mode.

Step 4 In ROMMOM mode, define the interface settings to the adaptive security appliance, including the IP address, TFTP server address, gateway address, software image file, and port, as follows:

```
rommon #1> ADDRESS=10.132.44.177
rommon #2> SERVER=10.129.0.30
rommon #3> GATEWAY=10.132.44.1
rommon #4> IMAGE=f1/asa800-232-k8.bin
rommon #5> PORT=Ethernet0/0
Ethernet0/0
Link is UP
MAC Address: 0012.d949.15b8
```

Note Be sure that the connection to the network already exists.

Step 5 To validate your settings, enter the **set** command:

```
rommon #6> set
ROMMON Variable Settings:
  ADDRESS=10.132.44.177
  SERVER=10.129.0.30
  GATEWAY=10.132.44.1
  PORT=Ethernet0/0
  VLAN=untagged
  IMAGE=f1/asa800-232-k8.bin
  CONFIG=
  LINKTIMEOUT=20
  PKTTIMEOUT=4
  RETRY=20
```

Step 6 Ping the TFTP server by entering the **ping server** command.

```
rommon #7> ping server
Sending 20, 100-byte ICMP Echoes to server 10.129.0.30, timeout is 4 seconds:

Success rate is 100 percent (20/20)
```

Step 7 Load the software image by entering the **tftp** command.

```
rommon #8> tftp
ROMMON Variable Settings:
  ADDRESS=10.132.44.177
  SERVER=10.129.0.30
  GATEWAY=10.132.44.1
  PORT=Ethernet0/0
  VLAN=untagged
  IMAGE=f1/asa800-232-k8.bin
  CONFIG=
  LINKTIMEOUT=20
  PKTTIMEOUT=4
  RETRY=20

tftp f1/asa800-232-k8.bin@10.129.0.30 via 10.132.44.1

Received 14450688 bytes

Launching TFTP Image...
Cisco PIX Security Appliance admin loader (3.0) #0: Mon Mar  5 16:00:07 MST 2007

Loading...
```

After the software image is successfully loaded, the adaptive security appliance automatically exits ROMMOM mode.

Step 8 To verify that the correct software image has been loaded into the adaptive security appliance, check the version in the adaptive security appliance by entering the following command:

```
hostname> show version
```

Erasing the Flash File System

Step 1 Connect to the adaptive security appliance console port.

Step 2 Power off the adaptive security appliance, and then power it on.

Step 3 During startup, press the **Escape** key when you are prompted to enter ROMMON mode.

Step 4 To erase the file system, enter the **erase** command, which overwrites all files and erases the file system, including hidden system files.

```
rommon #1> erase [disk0: | disk1: | flash:]
```

Other Troubleshooting Tools

The security appliance provides other troubleshooting tools that you can use. This section includes the following topics:

- Viewing Debug Messages, page A-12
- Capturing Packets, page A-12
- Viewing the Crash Dump, page A-12
- TACACS+ Server Lockout, page A-12
- Verifying that Server Authentication and Authorization are Working, page A-12
- User's Identity not Preserved Across Contexts, page A-13

Viewing Debug Messages

Because debugging output is assigned high priority in the CPU process, it can render the system unusable. For this reason, use **debug** commands only to troubleshoot specific problems or during troubleshooting sessions with Cisco TAC. Moreover, it is best to use **debug** commands during periods of less network traffic and fewer users. Debugging during these periods decreases the likelihood that increased debug command processing overhead will affect system use. To enable debug messages, see the **debug** commands in the *Cisco ASA 5500 Series Command Reference*.

Capturing Packets

Capturing packets is sometimes useful when troubleshooting connectivity problems or monitoring suspicious activity. We recommend contacting Cisco TAC if you want to use the packet capture feature. See the **capture** command in the *Cisco ASA 5500 Series Command Reference*.

Viewing the Crash Dump

If the security appliance crashes, you can view the crash dump information. We recommend contacting Cisco TAC if you want to interpret the crash dump. See the **show crashdump** command in the *Cisco ASA 5500 Series Command Reference*.

TACACS+ Server Lockout

We recommend that, when configuring TACACS+ server command authorization, you do not save your configuration until you are sure it works the way you expect. If you get locked out because of a mistake, you can usually recover access by restarting the security appliance. If you are still locked out, see Recovering from a Lockout, page 18-27.

Verifying that Server Authentication and Authorization are Working

To verify that the security appliance can contact an AAA server and authenticate or authorize a user, see Testing Server Authentication and Authorization, page 16-18.

User's Identity not Preserved Across Contexts

If your network will be organized into multiple contexts, be aware that, when changing contexts, the user identity is not preserved. The user becomes a default (enable_15) user in the new context, with Administrative access (privilege level 15 access).

Common Problems

This section describes common problems with the security appliance, and how you might resolve them.

Symptom ASDM window becomes blank when you click Configure.

Possible Cause CSDM failed due to the data.xml file.

Recommended Action Click **Refresh**.

Symptom The context configuration was not saved, and was lost when you reloaded.

Possible Cause You did not save each context within the context execution space. If you are configuring contexts at the command line, you did not save the current context before you changed to the next context.

Recommended Action Save each context within the context execution space using the **copy run start** command. You cannot save contexts from the system execution space.

Symptom You cannot make a Telnet or SSH connection to the security appliance interface.

Possible Cause You did not enable Telnet or SSH to the security appliance.

Recommended Action Enable Telnet or SSH to the security appliance.

Symptom You cannot ping the security appliance interface.

Possible Cause You disabled ICMP to the security appliance.

Recommended Action Enable ICMP to the security appliance for your IP address using the **icmp** command.

Symptom You cannot ping through the security appliance, although the access list allows it.

Possible Cause You did not enable the ICMP inspection engine or apply access lists on both the ingress and egress interfaces.

Recommended Action Because ICMP is a connectionless protocol, the security appliance does not automatically allow returning traffic through. In addition to an access list on the ingress interface, you either need to apply an access list to the egress interface to allow replying traffic, or enable the ICMP inspection engine, which treats ICMP connections as stateful connections.

Symptom Traffic does not pass between two interfaces on the same security level.

Possible Cause You did not enable the feature that allows traffic to pass between interfaces at the same security level.

Recommended Action Enable this feature.

APPENDIX B

Addresses, Protocols, and Ports

This appendix provides a quick reference for IP addresses, protocols, and applications. This appendix includes the following sections:

- IPv4 Addresses and Subnet Masks, page B-1
- IPv6 Addresses, page B-5
- Protocols and Applications, page B-11
- TCP and UDP Ports, page B-11
- Local Ports and Protocols, page B-14
- ICMP Types, page B-15

IPv4 Addresses and Subnet Masks

This section describes how to use IPv4 addresses in the security appliance. An IPv4 address is a 32-bit number written in dotted-decimal notation: four 8-bit fields (octets) converted from binary to decimal numbers, separated by dots. The first part of an IP address identifies the network on which the host resides, while the second part identifies the particular host on the given network. The network number field is called the network prefix. All hosts on a given network share the same network prefix but must have a unique host number. In classful IP, the class of the address determines the boundary between the network prefix and the host number.

This section includes the following topics:

- Classes, page B-1
- Private Networks, page B-2
- Subnet Masks, page B-2

Classes

IP host addresses are divided into three different address classes: Class A, Class B, and Class C. Each class fixes the boundary between the network prefix and the host number at a different point within the 32-bit address. Class D addresses are reserved for multicast IP.

- Class A addresses (1.xxx.xxx.xxx through 126.xxx.xxx.xxx) use only the first octet as the network prefix.

- Class B addresses (128.0.xxx.xxx through 191.255.xxx.xxx) use the first two octets as the network prefix.
- Class C addresses (192.0.0.xxx through 223.255.255.xxx) use the first three octets as the network prefix.

Because Class A addresses have 16,777,214 host addresses, and Class B addresses 65,534 hosts, you can use subnet masking to break these huge networks into smaller subnets.

Private Networks

If you need large numbers of addresses on your network, and they do not need to be routed on the Internet, you can use private IP addresses that the Internet Assigned Numbers Authority (IANA) recommends (see RFC 1918). The following address ranges are designated as private networks that should not be advertised:

- 10.0.0.0 through 10.255.255.255
- 172.16.0.0 through 172.31.255.255
- 192.168.0.0 through 192.168.255.255

Subnet Masks

A subnet mask lets you convert a single Class A, B, or C network into multiple networks. With a subnet mask, you can create an extended network prefix that adds bits from the host number to the network prefix. For example, a Class C network prefix always consists of the first three octets of the IP address. But a Class C extended network prefix uses part of the fourth octet as well.

Subnet masking is easy to understand if you use binary notation instead of dotted decimal. The bits in the subnet mask have a one-to-one correspondence with the Internet address:

- The bits are set to 1 if the corresponding bit in the IP address is part of the extended network prefix.
- The bits are set to 0 if the bit is part of the host number.

Example 1: If you have the Class B address 129.10.0.0 and you want to use the entire third octet as part of the extended network prefix instead of the host number, you must specify a subnet mask of 11111111.11111111.11111111.00000000. This subnet mask converts the Class B address into the equivalent of a Class C address, where the host number consists of the last octet only.

Example 2: If you want to use only part of the third octet for the extended network prefix, then you must specify a subnet mask like 11111111.11111111.11111000.00000000, which uses only 5 bits of the third octet for the extended network prefix.

You can write a subnet mask as a dotted-decimal mask or as a /*bits* ("slash *bits*") mask. In Example 1, for a dotted-decimal mask, you convert each binary octet into a decimal number: 255.255.255.0. For a /*bits* mask, you add the number of 1s: /24. In Example 2, the decimal number is 255.255.248.0 and the /bits is /21.

You can also supernet multiple Class C networks into a larger network by using part of the third octet for the extended network prefix. For example, 192.168.0.0/20.

This section includes the following topics:

- Determining the Subnet Mask, page B-3
- Determining the Address to Use with the Subnet Mask, page B-3

Determining the Subnet Mask

To determine the subnet mask based on how many hosts you want, see Table B-1.

Table B-1 Hosts, Bits, and Dotted-Decimal Masks

Hosts[1]	/Bits Mask	Dotted-Decimal Mask
16,777,216	/8	255.0.0.0 Class A Network
65,536	/16	255.255.0.0 Class B Network
32,768	/17	255.255.128.0
16,384	/18	255.255.192.0
8192	/19	255.255.224.0
4096	/20	255.255.240.0
2048	/21	255.255.248.0
1024	/22	255.255.252.0
512	/23	255.255.254.0
256	/24	255.255.255.0 Class C Network
128	/25	255.255.255.128
64	/26	255.255.255.192
32	/27	255.255.255.224
16	/28	255.255.255.240
8	/29	255.255.255.248
4	/30	255.255.255.252
Do not use	/31	255.255.255.254
1	/32	255.255.255.255 Single Host Address

[1] The first and last number of a subnet are reserved, except for /32, which identifies a single host.

Determining the Address to Use with the Subnet Mask

The following sections describe how to determine the network address to use with a subnet mask for a Class C-size and a Class B-size network. This section includes the following topics:

- Class C-Size Network Address, page B-3
- Class B-Size Network Address, page B-4

Class C-Size Network Address

For a network between 2 and 254 hosts, the fourth octet falls on a multiple of the number of host addresses, starting with 0. For example, the 8-host subnets (/29) of 192.168.0.x are as follows:

Subnet with Mask /29 (255.255.255.248)	Address Range[1]
192.168.0.0	192.168.0.0 to 192.168.0.7
192.168.0.8	192.168.0.8 to 192.168.0.15

Subnet with Mask /29 (255.255.255.248)	Address Range[1]
192.168.0.16	192.168.0.16 to 192.168.0.31
…	…
192.168.0.248	192.168.0.248 to 192.168.0.255

[1] The first and last address of a subnet are reserved. In the first subnet example, you cannot use 192.168.0.0 or 192.168.0.7.

Class B-Size Network Address

To determine the network address to use with the subnet mask for a network with between 254 and 65,534 hosts, you need to determine the value of the third octet for each possible extended network prefix. For example, you might want to subnet an address like 10.1.*x*.0, where the first two octets are fixed because they are used in the extended network prefix, and the fourth octet is 0 because all bits are used for the host number.

To determine the value of the third octet, follow these steps:

Step 1 Calculate how many subnets you can make from the network by dividing 65,536 (the total number of addresses using the third and fourth octet) by the number of host addresses you want.

For example, 65,536 divided by 4096 hosts equals 16.

Therefore, there are 16 subnets of 4096 addresses each in a Class B-size network.

Step 2 Determine the multiple of the third octet value by dividing 256 (the number of values for the third octet) by the number of subnets:

In this example, 256/16 = 16.

The third octet falls on a multiple of 16, starting with 0.

Therefore, the 16 subnets of the network 10.1 are as follows:

Subnet with Mask /20 (255.255.240.0)	Address Range[1]
10.1.0.0	10.1.0.0 to 10.1.15.255
10.1.16.0	10.1.16.0 to 10.1.31.255
10.1.32.0	10.1.32.0 to 10.1.47.255
…	…
10.1.240.0	10.1.240.0 to 10.1.255.255

[1] The first and last address of a subnet are reserved. In the first subnet example, you cannot use 10.1.0.0 or 10.1.15.255.

IPv6 Addresses

IPv6 is the next generation of the Internet Protocol after IPv4. It provides an expanded address space, a simplified header format, improved support for extensions and options, flow labeling capability, and authentication and privacy capabilities. IPv6 is described in RFC 2460. The IPv6 addressing architecture is described in RFC 3513.

This section describes the IPv6 address format and architecture and includes the following topics:

- IPv6 Address Format, page B-5
- IPv6 Address Types, page B-6
- IPv6 Address Prefixes, page B-10

Note This section describes the IPv6 address format, the types, and prefixes. For information about configuring the security appliance to use IPv6, see Chapter 9, "Configuring Interfaces."

IPv6 Address Format

IPv6 addresses are represented as a series of eight 16-bit hexadecimal fields separated by colons (:) in the format: x:x:x:x:x:x:x:x. The following are two examples of IPv6 addresses:

- 2001:0DB8:7654:3210:FEDC:BA98:7654:3210
- 2001:0DB8:0000:0000:0008:0800:200C:417A

Note The hexadecimal letters in IPv6 addresses are not case-sensitive.

It is not necessary to include the leading zeros in an individual field of the address. But each field must contain at least one digit. So the example address 2001:0DB8:0000:0000:0008:0800:200C:417A can be shortened to 2001:0DB8:0:0:8:800:200C:417A by removing the leading zeros from the third through sixth fields from the left. The fields that contained all zeros (the third and fourth fields from the left) were shortened to a single zero. The fifth field from the left had the three leading zeros removed, leaving a single 8 in that field, and the sixth field from the left had the one leading zero removed, leaving 800 in that field.

It is common for IPv6 addresses to contain several consecutive hexadecimal fields of zeros. You can use two colons (::) to compress consecutive fields of zeros at the beginning, middle, or end of an IPv6 address (the colons represent the successive hexadecimal fields of zeros). Table B-2 shows several examples of address compression for different types of IPv6 address.

Table B-2 IPv6 Address Compression Examples

Address Type	Standard Form	Compressed Form
Unicast	2001:0DB8:0:0:0:BA98:0:3210	2001:0DB8::BA98:0:3210
Multicast	FF01:0:0:0:0:0:0:101	FF01::101
Loopback	0:0:0:0:0:0:0:1	::1
Unspecified	0:0:0:0:0:0:0:0	::

Note Two colons (::) can be used only once in an IPv6 address to represent successive fields of zeros.

An alternative form of the IPv6 format is often used when dealing with an environment that contains both IPv4 and IPv6 addresses. This alternative has the format x:x:x:x:x:x:y.y.y.y, where x represent the hexadecimal values for the six high-order parts of the IPv6 address and y represent decimal values for the 32-bit IPv4 part of the address (which takes the place of the remaining two 16-bit parts of the IPv6 address). For example, the IPv4 address 192.168.1.1 could be represented as the IPv6 address 0:0:0:0:0:0:FFFF:192.168.1.1, or ::FFFF:192.168.1.1.

IPv6 Address Types

The following are the three main types of IPv6 addresses:

- **Unicast**—A unicast address is an identifier for a single interface. A packet sent to a unicast address is delivered to the interface identified by that address. An interface may have more than one unicast address assigned to it.
- **Multicast**—A multicast address is an identifier for a set of interfaces. A packet sent to a multicast address is delivered to all addresses identified by that address.
- **Anycast**—An anycast address is an identifier for a set of interfaces. Unlike a multicast address, a packet sent to an anycast address is only delivered to the “nearest” interface, as determined by the measure of distances for the routing protocol.

Note There are no broadcast addresses in IPv6. Multicast addresses provide the broadcast functionality.

This section includes the following topics:

- Unicast Addresses, page B-6
- Multicast Address, page B-8
- Anycast Address, page B-9
- Required Addresses, page B-10

Unicast Addresses

This section describes IPv6 unicast addresses. Unicast addresses identify an interface on a network node.

This section includes the following topics:

- Global Address, page B-7
- Site-Local Address, page B-7
- Link-Local Address, page B-7
- IPv4-Compatible IPv6 Addresses, page B-7
- Unspecified Address, page B-8
- Loopback Address, page B-8
- Interface Identifiers, page B-8

Global Address

The general format of an IPv6 global unicast address is a global routing prefix followed by a subnet ID followed by an interface ID. The global routing prefix can be any prefix not reserved by another IPv6 address type (see IPv6 Address Prefixes, page B-10, for information about the IPv6 address type prefixes).

All global unicast addresses, other than those that start with binary 000, have a 64-bit interface ID in the Modified EUI-64 format. See Interface Identifiers, page B-8, for more information about the Modified EUI-64 format for interface identifiers.

Global unicast address that start with the binary 000 do not have any constraints on the size or structure of the interface ID portion of the address. One example of this type of address is an IPv6 address with an embedded IPv4 address (see IPv4-Compatible IPv6 Addresses, page B-7).

Site-Local Address

Site-local addresses are used for addressing within a site. They can be use to address an entire site without using a globally unique prefix. Site-local addresses have the prefix FEC0::/10, followed by a 54-bit subnet ID, and end with a 64-bit interface ID in the modified EUI-64 format.

Site-local Routers do not forward any packets that have a site-local address for a source or destination outside of the site. Therefore, site-local addresses can be considered private addresses.

Link-Local Address

All interfaces are required to have at least one link-local address. You can configure multiple IPv6 addresses per interfaces, but only one link-local address.

A link-local address is an IPv6 unicast address that can be automatically configured on any interface using the link-local prefix FE80::/10 and the interface identifier in modified EUI-64 format. Link-local addresses are used in the neighbor discovery protocol and the stateless autoconfiguration process. Nodes with a link-local address can communicate; they do not need a site-local or globally unique address to communicate.

Routers do not forward any packets that have a link-local address for a source or destination. Therefore, link-local addresses can be considered private addresses.

IPv4-Compatible IPv6 Addresses

There are two types of IPv6 addresses that can contain IPv4 addresses.

The first type is the “IPv4-compatibly IPv6 address.” The IPv6 transition mechanisms include a technique for hosts and routers to dynamically tunnel IPv6 packets over IPv4 routing infrastructure. IPv6 nodes that use this technique are assigned special IPv6 unicast addresses that carry a global IPv4 address in the low-order 32 bits. This type of address is termed an “IPv4-compatible IPv6 address” and has the format ::y.y.y.y, where y.y.y.y is an IPv4 unicast address.

Note The IPv4 address used in the “IPv4-compatible IPv6 address” must be a globally-unique IPv4 unicast address.

The second type of IPv6 address which holds an embedded IPv4 address is called the “IPv4-mapped IPv6 address.” This address type is used to represent the addresses of IPv4 nodes as IPv6 addresses. This type of address has the format ::FFFF:y.y.y.y, where y.y.y.y is an IPv4 unicast address.

Unspecified Address

The unspecified address, 0:0:0:0:0:0:0:0, indicates the absence of an IPv6 address. For example, a newly initialized node on an IPv6 network may use the unspecified address as the source address in its packets until it receives its IPv6 address.

Note The IPv6 unspecified address cannot be assigned to an interface. The unspecified IPv6 addresses must not be used as destination addresses in IPv6 packets or the IPv6 routing header.

Loopback Address

The loopback address, 0:0:0:0:0:0:0:1, may be used by a node to send an IPv6 packet to itself. The loopback address in IPv6 functions the same as the loopback address in IPv4 (127.0.0.1).

Note The IPv6 loopback address cannot be assigned to a physical interface. A packet that has the IPv6 loopback address as its source or destination address must remain within the node that created the packet. IPv6 routers do not forward packets that have the IPv6 loopback address as their source or destination address.

Interface Identifiers

Interface identifiers in IPv6 unicast addresses are used to identify the interfaces on a link. They need to be unique within a subnet prefix. In many cases, the interface identifier is derived from the interface link-layer address. The same interface identifier may be used on multiple interfaces of a single node, as long as those interfaces are attached to different subnets.

For all unicast addresses, except those that start with the binary 000, the interface identifier is required to be 64 bits long and to be constructed in the Modified EUI-64 format. The Modified EUI-64 format is created from the 48-bit MAC address by inverting the universal/local bit in the address and by inserting the hexadecimal number FFFE between the upper three bytes and lower three bytes of the of the MAC address.

For example, and interface with the MAC address of 00E0.b601.3B7A would have a 64-bit interface ID of 02E0:B6FF:FE01:3B7A.

Multicast Address

An IPv6 multicast address is an identifier for a group of interfaces, typically on different nodes. A packet sent to a multicast address is delivered to all interfaces identified by the multicast address. An interface may belong to any number of multicast groups.

An IPv6 multicast address has a prefix of FF00::/8 (1111 1111). The octet following the prefix defines the type and scope of the multicast address. A permanently assigned (“well known”) multicast address has a flag parameter equal to 0; a temporary (“transient”) multicast address has a flag parameter equal to 1. A multicast address that has the scope of a node, link, site, or organization, or a global scope has a scope parameter of 1, 2, 5, 8, or E, respectively. For example, a multicast address with the prefix FF02::/16 is a permanent multicast address with a link scope. Figure B-1 shows the format of the IPv6 multicast address.

Figure B-1 IPv6 Multicast Address Format

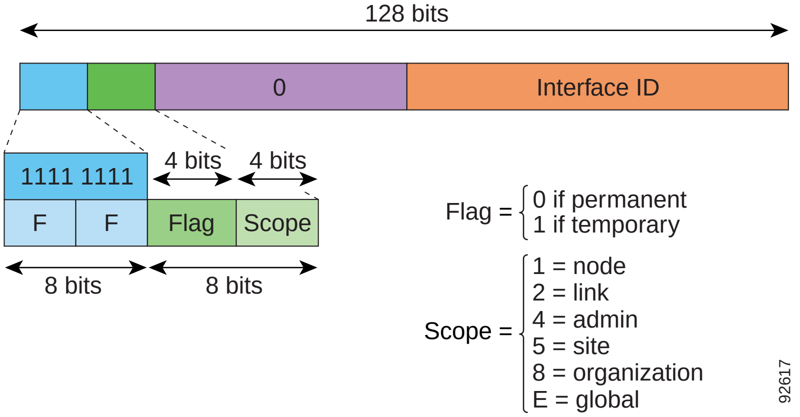

IPv6 nodes (hosts and routers) are required to join the following multicast groups:

- The All Nodes multicast addresses:
  - FF01:: (interface-local)
  - FF02:: (link-local)
- The Solicited-Node Address for each IPv6 unicast and anycast address on the node: FF02:0:0:0:0:1:FFXX:XXXX/104, where XX:XXXX is the low-order 24-bits of the unicast or anycast address.

Note Solicited-Node addresses are used in Neighbor Solicitation messages.

IPv6 routers are required to join the following multicast groups:

- FF01::2 (interface-local)
- FF02::2 (link-local)
- FF05::2 (site-local)

Multicast address should not be used as source addresses in IPv6 packets.

Note There are no broadcast addresses in IPv6. IPv6 multicast addresses are used instead of broadcast addresses.

Anycast Address

The IPv6 anycast address is a unicast address that is assigned to more than one interface (typically belonging to different nodes). A packet that is routed to an anycast address is routed to the nearest interface having that address, the nearness being determined by the routing protocol in effect.

Anycast addresses are allocated from the unicast address space. An anycast address is simply a unicast address that has been assigned to more than one interface, and the interfaces must be configured to recognize the address as an anycast address.

The following restrictions apply to anycast addresses:

- An anycast address cannot be used as the source address for an IPv6 packet.

- An anycast address cannot be assigned to an IPv6 host; it can only be assigned to an IPv6 router.

Note Anycast addresses are not supported on the security appliance.

Required Addresses

IPv6 hosts must, at a minimum, be configured with the following addresses (either automatically or manually):

- A link-local address for each interface.
- The loopback address.
- The All-Nodes multicast addresses
- A Solicited-Node multicast address for each unicast or anycast address.

IPv6 routers must, at a minimum, be configured with the following addresses (either automatically or manually):

- The required host addresses.
- The Subnet-Router anycast addresses for all interfaces for which it is configured to act as a router.
- The All-Routers multicast addresses.

IPv6 Address Prefixes

An IPv6 address prefix, in the format ipv6-prefix/prefix-length, can be used to represent bit-wise contiguous blocks of the entire address space. The IPv6-prefix must be in the form documented in RFC 2373 where the address is specified in hexadecimal using 16-bit values between colons. The prefix length is a decimal value that indicates how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address). For example, 2001:0DB8:8086:6502::/32 is a valid IPv6 prefix.

The IPv6 prefix identifies the type of IPv6 address. Table B-3 shows the prefixes for each IPv6 address type.

Table B-3 IPv6 Address Type Prefixes

Address Type	Binary Prefix	IPv6 Notation
Unspecified	000...0 (128 bits)	::/128
Loopback	000...1 (128 bits)	::1/128
Multicast	11111111	FF00::/8
Link-Local (unicast)	1111111010	FE80::/10
Site-Local (unicast)	1111111111	FEC0::/10
Global (unicast)	All other addresses.	
Anycast	Taken from the unicast address space.	

Protocols and Applications

Table B-4 lists the protocol literal values and port numbers; either can be entered in security appliance commands.

Table B-4 Protocol Literal Values

Literal	Value	Description
ah	51	Authentication Header for IPv6, RFC 1826.
eigrp	88	Enhanced Interior Gateway Routing Protocol.
esp	50	Encapsulated Security Payload for IPv6, RFC 1827.
gre	47	Generic Routing Encapsulation.
icmp	1	Internet Control Message Protocol, RFC 792.
icmp6	58	Internet Control Message Protocol for IPv6, RFC 2463.
igmp	2	Internet Group Management Protocol, RFC 1112.
igrp	9	Interior Gateway Routing Protocol.
ip	0	Internet Protocol.
ipinip	4	IP-in-IP encapsulation.
ipsec	50	IP Security. Entering the ipsec protocol literal is equivalent to entering the esp protocol literal.
nos	94	Network Operating System (Novell's NetWare).
ospf	89	Open Shortest Path First routing protocol, RFC 1247.
pcp	108	Payload Compression Protocol.
pim	103	Protocol Independent Multicast.
pptp	47	Point-to-Point Tunneling Protocol. Entering the pptp protocol literal is equivalent to entering the gre protocol literal.
snp	109	Sitara Networks Protocol.
tcp	6	Transmission Control Protocol, RFC 793.
udp	17	User Datagram Protocol, RFC 768.

Protocol numbers can be viewed online at the IANA website:

http://www.iana.org/assignments/protocol-numbers

TCP and UDP Ports

Table B-5 lists the literal values and port numbers; either can be entered in security appliance commands. See the following caveats:

- The security appliance uses port 1521 for SQL*Net. This is the default port used by Oracle for SQL*Net. This value, however, does not agree with IANA port assignments.

- The security appliance listens for RADIUS on ports 1645 and 1646. If your RADIUS server uses the standard ports 1812 and 1813, you can configure the security appliance to listen to those ports using the **authentication-port** and **accounting-port** commands.
- To assign a port for DNS access, use the **domain** literal value, not **dns**. If you use **dns**, the security appliance assumes you meant to use the **dnsix** literal value.

Port numbers can be viewed online at the IANA website:

http://www.iana.org/assignments/port-numbers

Table B-5 Port Literal Values

Literal	TCP or UDP?	Value	Description
aol	TCP	5190	America Online
bgp	TCP	179	Border Gateway Protocol, RFC 1163
biff	UDP	512	Used by mail system to notify users that new mail is received
bootpc	UDP	68	Bootstrap Protocol Client
bootps	UDP	67	Bootstrap Protocol Server
chargen	TCP	19	Character Generator
citrix-ica	TCP	1494	Citrix Independent Computing Architecture (ICA) protocol
cmd	TCP	514	Similar to **exec** except that **cmd** has automatic authentication
ctiqbe	TCP	2748	Computer Telephony Interface Quick Buffer Encoding
daytime	TCP	13	Day time, RFC 867
discard	TCP, UDP	9	Discard
domain	TCP, UDP	53	DNS
dnsix	UDP	195	DNSIX Session Management Module Audit Redirector
echo	TCP, UDP	7	Echo
exec	TCP	512	Remote process execution
finger	TCP	79	Finger
ftp	TCP	21	File Transfer Protocol (control port)
ftp-data	TCP	20	File Transfer Protocol (data port)
gopher	TCP	70	Gopher
https	TCP	443	HTTP over SSL
h323	TCP	1720	H.323 call signalling
hostname	TCP	101	NIC Host Name Server
ident	TCP	113	Ident authentication service
imap4	TCP	143	Internet Message Access Protocol, version 4
irc	TCP	194	Internet Relay Chat protocol

Table B-5 Port Literal Values (continued)

Literal	TCP or UDP?	Value	Description
isakmp	UDP	500	Internet Security Association and Key Management Protocol
kerberos	TCP, UDP	750	Kerberos
klogin	TCP	543	KLOGIN
kshell	TCP	544	Korn Shell
ldap	TCP	389	Lightweight Directory Access Protocol
ldaps	TCP	636	Lightweight Directory Access Protocol (SSL)
lpd	TCP	515	Line Printer Daemon - printer spooler
login	TCP	513	Remote login
lotusnotes	TCP	1352	IBM Lotus Notes
mobile-ip	UDP	434	MobileIP-Agent
nameserver	UDP	42	Host Name Server
netbios-ns	UDP	137	NetBIOS Name Service
netbios-dgm	UDP	138	NetBIOS Datagram Service
netbios-ssn	TCP	139	NetBIOS Session Service
nntp	TCP	119	Network News Transfer Protocol
ntp	UDP	123	Network Time Protocol
pcanywhere-status	UDP	5632	pcAnywhere status
pcanywhere-data	TCP	5631	pcAnywhere data
pim-auto-rp	TCP, UDP	496	Protocol Independent Multicast, reverse path flooding, dense mode
pop2	TCP	109	Post Office Protocol - Version 2
pop3	TCP	110	Post Office Protocol - Version 3
pptp	TCP	1723	Point-to-Point Tunneling Protocol
radius	UDP	1645	Remote Authentication Dial-In User Service
radius-acct	UDP	1646	Remote Authentication Dial-In User Service (accounting)
rip	UDP	520	Routing Information Protocol
secureid-udp	UDP	5510	SecureID over UDP
smtp	TCP	25	Simple Mail Transport Protocol
snmp	UDP	161	Simple Network Management Protocol
snmptrap	UDP	162	Simple Network Management Protocol - Trap
sqlnet	TCP	1521	Structured Query Language Network
ssh	TCP	22	Secure Shell
sunrpc (rpc)	TCP, UDP	111	Sun Remote Procedure Call
syslog	UDP	514	System Log

Table B-5 Port Literal Values (continued)

Literal	TCP or UDP?	Value	Description
tacacs	TCP, UDP	49	Terminal Access Controller Access Control System Plus
talk	TCP, UDP	517	Talk
telnet	TCP	23	RFC 854 Telnet
tftp	UDP	69	Trivial File Transfer Protocol
time	UDP	37	Time
uucp	TCP	540	UNIX-to-UNIX Copy Program
who	UDP	513	Who
whois	TCP	43	Who Is
www	TCP	80	World Wide Web
xdmcp	UDP	177	X Display Manager Control Protocol

Local Ports and Protocols

Table B-6 lists the protocols, TCP ports, and UDP ports that the security appliance may open to process traffic destined to the security appliance. Unless you enable the features and services listed in Table B-6, the security appliance does *not* open any local protocols or any TCP or UDP ports. You must configure a feature or service for the security appliance to open the default listening protocol or port. In many cases you can configure ports other than the default port when you enable a feature or service.

Table B-6 Protocols and Ports Opened by Features and Services

Feature or Service	Protocol	Port Number	Comments
DHCP	UDP	67,68	—
Failover Control	108	N/A	—
HTTP	TCP	80	—
HTTPS	TCP	443	—
ICMP	1	N/A	—
IGMP	2	N/A	Protocol only open on destination IP address 224.0.0.1
ISAKMP/IKE	UDP	500	Configurable.
IPSec (ESP)	50	N/A	—
IPSec over UDP (NAT-T)	UDP	4500	—
IPSec over UDP (Cisco VPN 3000 Series compatible)	UDP	10000	Configurable.
IPSec over TCP (CTCP)	TCP	—	No default port is used. You must specify the port number when configuring IPSec over TCP.

Table B-6 Protocols and Ports Opened by Features and Services (continued)

Feature or Service	Protocol	Port Number	Comments
NTP	UDP	123	—
OSPF	89	N/A	Protocol only open on destination IP address 224.0.0.5 and 224.0.0.6
PIM	103	N/A	Protocol only open on destination IP address 224.0.0.13
RIP	UDP	520	—
RIPv2	UDP	520	Port only open on destination IP address 224.0.0.9
SNMP	UDP	161	Configurable.
SSH	TCP	22	—
Stateful Update	105	N/A	—
Telnet	TCP	23	—
VPN Load Balancing	UDP	9023	Configurable.
VPN Individual User Authentication Proxy	UDP	1645, 1646	Port accessible only over VPN tunnel.

ICMP Types

Table B-7 lists the ICMP type numbers and names that you can enter in security appliance commands:

Table B-7 ICMP Types

ICMP Number	ICMP Name
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

Table B-7 ICMP Types (continued)

ICMP Number	ICMP Name
18	mask-reply
31	conversion-error
32	mobile-redirect

APPENDIX C

Configuring an External Server for Authorization and Authentication

This appendix describes how to configure an external LDAP, RADIUS, or TACACS+ server to support AAA on the security appliance. Before you configure the security appliance to use an external server, you must configure the server with the correct security appliance authorization attributes and, from a subset of these attributes, assign specific permissions to individual users.

This appendix includes the following sections:

Understanding Policy Enforcement of Permissions and Attributes

The security appliance supports several methods of applying user authorization attributes (also called user entitlements or permissions) to VPN connections. You can configure the security appliance to obtain user attributes from a Dynamic Access Policy (DAP) on the security appliance, from an external authentication and/or authorization AAA server (RADIUS or LDAP), from a group policy on the security appliance, or from all three.

If the security appliance receives attributes from all sources, the attributes are evaluated, merged, and applied to the user policy. If there are conflicts between attributes coming from the DAP, the AAA server, or the group policy, those attributes obtained from the DAP always take precedence.

The security appliance applies attributes in the following order (also illustrated in Figure C-1:

1. DAP attributes on the security appliance—Introduced in Version 8.0, take precedence over all others. If you set a bookmark/URL list in DAP, it overrides a bookmark/URL list set in the group policy.
2. User attributes on the AAA server—The server returns these after successful user authentication and/or authorization. Do not confuse these with attributes that are set for individual users in the local AAA database on the security appliance (User Accounts in ASDM).
3. Group policy configured on the security appliance—If a RADIUS server returns the value of the RADIUS CLASS attribute IETF-Class-25 (OU=<group-policy>) for the user, the security appliance places the user in the group policy of the same name and enforces any attributes in the group policy that are not returned by the server.

   For LDAP servers, any attribute name can be used to set the group policy for the session. The LDAP attribute map you configure on the security appliance maps the LDAP attribute to the Cisco attribute IETF-Radius-Class.
4. Group policy assigned by the Connection Profile (called tunnel-group in CLI)—The Connection Profile has the preliminary settings for the connection, and includes a default group policy applied to the user before authentication. All users connecting to the security appliance initially belong to this group which provides any attributes that are missing from the DAP, user attributes returned by the server, or the group policy assigned to the user.
5. Default group policy assigned by the security appliance (DfltGrpPolicy)—System default attributes provide any values that are missing from the DAP, user attributes, group policy, or connection profile.

Figure C-1 Policy Enforcement Flow

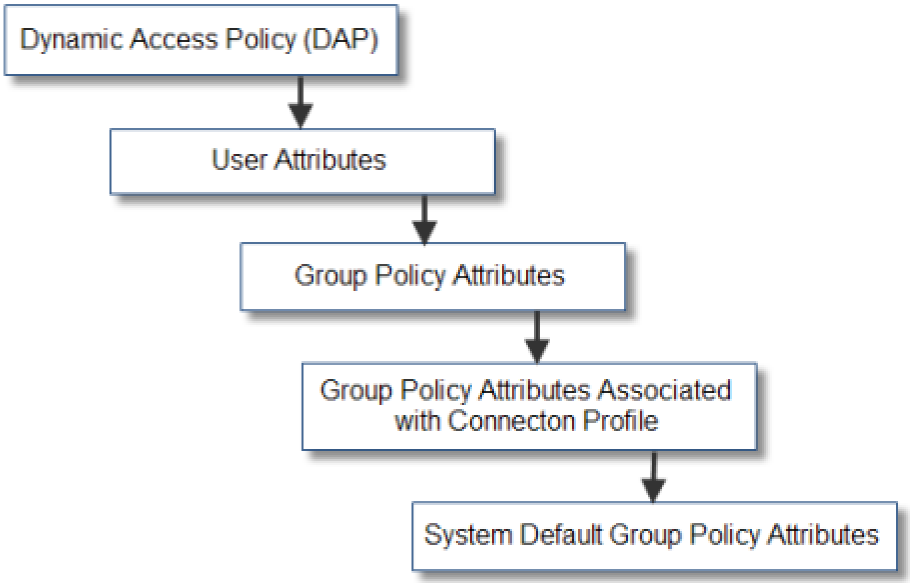
Dynamic Access Policy (DAP)
User Attributes
Group Policy Attributes
Group Policy Attributes Associated with Connecton Profile
System Default Group Policy Attributes

Configuring an External LDAP Server

The VPN 3000 Concentrator and the ASA/PIX 7.0 required a Cisco LDAP schema for authorization operations. Beginning with Version 7.1.x, the security appliance performs authentication *and* authorization, using the native LDAP schema, and the Cisco schema is no longer needed.

You configure authorization (permission policy) using an LDAP attribute map. For examples, see Active Directory/LDAP VPN Remote Access Authorization Use Cases, page C-16.

This section describes the structure, schema, and attributes of an LDAP server. It includes the following topics:

- Organizing the Security Appliance for LDAP Operations, page C-3
- Defining the Security Appliance LDAP Configuration, page C-6
- Active Directory/LDAP VPN Remote Access Authorization Use Cases, page C-16

The specific steps of these processes vary, depending on which type of LDAP server you are using.

Note For more information on the LDAP protocol, see RFCs 1777, 2251, and 2849.

Organizing the Security Appliance for LDAP Operations

This section describes how to perform searches within the LDAP hierarchy and authenticated binding to the LDAP server on the security appliance. It includes the following topics:

- Searching the Hierarchy, page C-4
- Binding the Security Appliance to the LDAP Server, page C-5
- Login DN Example for Active Directory, page C-5

Your LDAP configuration should reflect the logical hierarchy of your organization. For example, suppose an employee at your company, Example Corporation, is named Terry. Terry works in the Engineering group. Your LDAP hierarchy could have one or many levels. You might decide to set up a shallow, single-level hierarchy in which Terry is considered a member of Example Corporation. Or, you could set up a multi-level hierarchy in which Terry is considered to be a member of the department Engineering, which is a member of an organizational unit called People, which is itself a member of Example Corporation. See Figure C-2 for an example of this multi-level hierarchy.

A multi-level hierarchy has more granularity, but a single level hierarchy is quicker to search.

Figure C-2 A Multi-Level LDAP Hierarchy

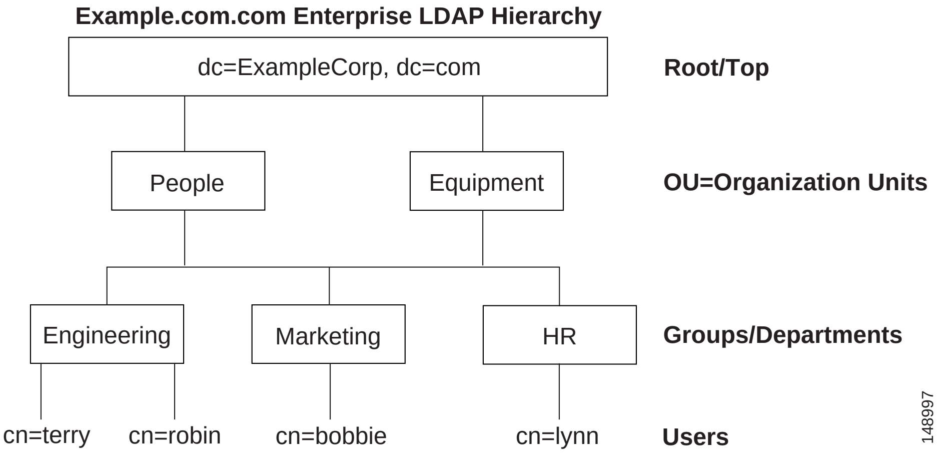

Searching the Hierarchy

The security appliance lets you tailor the search within the LDAP hierarchy. You configure the following three fields on the security appliance to define where in the LDAP hierarchy your search begins, the extent, and the type of information it is looking for. Together these fields allow you to limit the search of the hierarchy to only the part of the tree that contains the user permissions.

- LDAP Base DN defines where in the LDAP hierarchy the server should begin searching for user information when it receives an authorization request from the security appliance.
- Search Scope defines the extent of the search in the LDAP hierarchy. The search proceeds this many levels in the hierarchy below the LDAP Base DN. You can choose to have the server search only the level immediately below, or it can search the entire subtree. A single level search is quicker, but a subtree search is more extensive.
- Naming Attribute(s) defines the RDN that uniquely identifies an entry in the LDAP server. Common naming attributes can include cn (Common Name), sAMAccountName, and userPrincipalName.

Figure C-2 shows a possible LDAP hierarchy for Example Corporation. Given this hierarchy, you could define your search in different ways. Table C-1 shows two possible search configurations.

In the first example configuration, when Terry establishes the IPSec tunnel with LDAP authorization required, the security appliance sends a search request to the LDAP server indicating it should search for Terry in the Engineering group. This search is quick.

In the second example configuration, the security appliance sends a search request indicating the server should search for Terry within Example Corporation. This search takes longer.

Table C-1 Example Search Configurations

#	LDAP Base DN	Search Scope	Naming Attribute	Result
1	group= Engineering,ou=People,dc=ExampleCorporation, dc=com	One Level	cn=Terry	Quicker search
2	dc=ExampleCorporation,dc=com	Subtree	cn=Terry	Longer search

Binding the Security Appliance to the LDAP Server

Some LDAP servers (including the Microsoft Active Directory server) require the security appliance to establish a handshake via authenticated binding before they accept requests for any other LDAP operations. The security appliance uses the Login Distinguished Name (DN) and Login Password to establish trust (bind) with an LDAP server. The Login DN represents a user record in the LDAP server that the administrator uses for binding.

When binding, the security appliance authenticates to the server using the Login DN and the Login Password. When performing a Microsoft Active Directory read-only operation (such as for authentication, authorization, or group-search), the security appliance can bind with a Login DN with less privileges. For example, the Login DN can be a user whose AD "Member Of" designation is part of Domain Users. For VPN password management operations, the Login DN needs elevated privileges and must be part of the Account Operators AD group.

An example of a Login DN includes:

cn=Binduser1,ou=Admins,ou=Users,dc=company_A,dc=com

The security appliance supports:

- Simple LDAP authentication with an unencrypted password on port 389
- Secure LDAP (LDAP-S) on port 636
- Simple Authentication and Security Layer (SASL) MD5
- SASL Kerberos.

The security appliance does not support anonymous authentication.

Note As an LDAP client, the security appliance does not support sending anonymous binds or requests.

Login DN Example for Active Directory

The Login DN is a username on the LDAP server that the security appliance uses to establish a trust between itself (the LDAP client) and the LDAP server during the Bind exchange, before a user search can take place.

For VPN authentication/authorization operations, and beginning with version 8.0.4 for retrieval of AD Groups, (which are read operations only when password-management changes are not required), the you can use the Login DN with fewer privileges. For example, the Login DN can be a user who is a memberOf the Domain Users group.

For VPN password-management changes, the Login DN must have Account Operators privileges.

In either of these cases, Super-user level privileges are not required for the Login/Bind DN. Refer to your LDAP Administrator guide for specific Login DN requirements.

Defining the Security Appliance LDAP Configuration

This section describes how to define the LDAP AV-pair attribute syntax. It includes the following topics:

- Supported Cisco Attributes for LDAP Authorization, page C-6
- Cisco AV Pair Attribute Syntax, page C-13
- Cisco AV Pairs ACL Examples, page C-15

Note The security appliance enforces the LDAP attributes based on attribute name, not numeric ID. RADIUS attributes, on the other hand, are enforced by numeric ID, not by name.

Authorization refers to the process of enforcing permissions or attributes. An LDAP server defined as an authentication or authorization server will enforce permissions or attributes if they are configured.

For software Version 7.0, LDAP attributes include the cVPN3000 prefix. For Version 7.1 and later, this prefix was removed.

Supported Cisco Attributes for LDAP Authorization

This section provides a complete list of attributes (Table C-2) for the ASA 5500, VPN 3000, and PIX 500 series security appliances. The table includes attribute support information for the VPN 3000 and PIX 500 series to assist you configure networks with a mixture of these security appliances.

Table C-2 Security Appliance Supported Cisco Attributes for LDAP Authorization

Attribute Name/	VPN 3000	ASA	PIX	Syntax/ Type	Single or Multi-Valued	Possible Values
Access-Hours	Y	Y	Y	String	Single	Name of the time-range (for example, Business-Hours)
Allow-Network-Extension- Mode	Y	Y	Y	Boolean	Single	0 = Disabled 1 = Enabled
Authenticated-User-Idle- Timeout	Y	Y	Y	Integer	Single	1 - 35791394 minutes
Authorization-Required	Y			Integer	Single	0 = No 1 = Yes
Authorization-Type	Y			Integer	Single	0 = None 1 = RADIUS 2 = LDAP
Banner1	Y	Y	Y	String	Single	Banner string for clientless and client SSL VPN, and IPSec clients.
Banner2	Y	Y	Y	String	Single	Banner string for clientless and client SSL VPN, and IPSec clients.

Table C-2 Security Appliance Supported Cisco Attributes for LDAP Authorization (continued)

Attribute Name/	VPN 3000	ASA	PIX	Syntax/ Type	Single or Multi-Valued	Possible Values
Cisco-AV-Pair	Y	Y	Y	String	Multi	An octet string in the following format: [Prefix] [Action] [Protocol] [Source] [Source Wildcard Mask] [Destination] [Destination Wildcard Mask] [Established] [Log] [Operator] [Port] For more information, see "Cisco AV Pair Attribute Syntax."
Cisco-IP-Phone-Bypass	Y	Y	Y	Integer	Single	0 = Disabled 1 = Enabled
Cisco-LEAP-Bypass	Y	Y	Y	Integer	Single	0 = Disabled 1 = Enabled
Client-Intercept-DHCP-Configure-Msg	Y	Y	Y	Boolean	Single	0 = Disabled 1 = Enabled
Client-Type-Version-Limiting	Y	Y	Y	String	Single	IPSec VPN client version number string
Confidence-Interval	Y	Y	Y	Integer	Single	10 - 300 seconds
DHCP-Network-Scope	Y	Y	Y	String	Single	IP address
DN-Field	Y	Y	Y	String	Single	Possible values: UID, OU, O, CN, L, SP, C, EA, T, N, GN, SN, I, GENQ, DNQ, SER, use-entire-name.
Firewall-ACL-In		Y	Y	String	Single	Access list ID
Firewall-ACL-Out		Y	Y	String	Single	Access list ID
Group-Policy		Y	Y	String	Single	Sets the group policy for the remote access VPN session. For version 8.2 and later, use this attribute instead of IETF-Radius-Class. You can use one of the three following formats: • <group policy name> • OU=<group policy name> • OU=<group policy name>;
IE-Proxy-Bypass-Local				Boolean	Single	0=Disabled 1=Enabled
IE-Proxy-Exception-List				String	Single	A list of DNS domains. Entries must be separated by the new line character sequence (\n).

Table C-2 Security Appliance Supported Cisco Attributes for LDAP Authorization (continued)

Attribute Name/	VPN 3000	ASA	PIX	Syntax/ Type	Single or Multi-Valued	Possible Values
IE-Proxy-Method	Y	Y	Y	Integer	Single	1 = Do not modify proxy settings 2 = Do not use proxy 3 = Auto detect 4 = Use security appliance setting
IE-Proxy-Server	Y	Y	Y	Integer	Single	IP Address
IETF-Radius-Class	Y	Y	Y		Single	Sets the group policy for the remote access VPN session. For version 8.2 and later, we recommend that you use the Group-Policy attribute. You can use one of the three following formats: • <group policy name> • OU=<group policy name> • OU=<group policy name>;
IETF-Radius-Filter-Id	Y	Y	Y	String	Single	access list name that is defined on the security appliance
IETF-Radius-Framed-IP-Address	Y	Y	Y	String	Single	An IP address
IETF-Radius-Framed-IP-Netmask	Y	Y	Y	String	Single	An IP address mask
IETF-Radius-Idle-Timeout	Y	Y	Y	Integer	Single	seconds
IETF-Radius-Service-Type	Y	Y	Y	Integer	Single	1 = Login 2 = Framed 6 = Administrative 7 = NAS Prompt
IETF-Radius-Session-Timeout	Y	Y	Y	Integer	Single	seconds
IKE-Keep-Alives	Y	Y	Y	Boolean	Single	0 = Disabled 1 = Enabled
IPSec-Allow-Passwd-Store	Y	Y	Y	Boolean	Single	0 = Disabled 1 = Enabled
IPSec-Authentication	Y	Y	Y	Integer	Single	0 = None 1 = RADIUS 2 = LDAP (authorization only) 3 = NT Domain 4 = SDI (RSA) 5 = Internal 6 = RADIUS with Expiry 7 = Kerberos/Active Directory
IPSec-Auth-On-Rekey	Y	Y	Y	Boolean	Single	0 = Disabled 1 = Enabled
IPSec-Backup-Server-List	Y	Y	Y	String	Single	Server Addresses (space delimited)
IPSec-Backup-Servers	Y	Y	Y	String	Single	1 = Use Client-Configured list 2 = Disabled and clear client list 3 = Use Backup Server list

Table C-2 Security Appliance Supported Cisco Attributes for LDAP Authorization (continued)

Attribute Name/	VPN 3000	ASA	PIX	Syntax/ Type	Single or Multi-Valued	Possible Values
IPSec-Client-Firewall-Filter- Name	Y			String	Single	Specifies the name of the filter to be pushed to the client as firewall policy.
IPSec-Client-Firewall-Filter-Optional	Y	Y	Y	Integer	Single	0 = Required 1 = Optional
IPSec-Default-Domain	Y	Y	Y	String	Single	Specifies the single default domain name to send to the client (1 - 255 characters).
IPSec-Extended-Auth-On-Rekey		Y	Y	String	Single	
IPSec-IKE-Peer-ID-Check	Y	Y	Y	Integer	Single	1 = Required 2 = If supported by peer certificate 3 = Do not check
IPSec-IP-Compression	Y	Y	Y	Integer	Single	0 = Disabled 1 = Enabled
IPSec-Mode-Config	Y	Y	Y	Boolean	Single	0 = Disabled 1 = Enabled
IPSec-Over-UDP	Y	Y	Y	Boolean	Single	0 = Disabled 1 = Enabled
IPSec-Over-UDP-Port	Y	Y	Y	Integer	Single	4001 - 49151; default = 10000
IPSec-Required-Client-Firewall-Capability	Y	Y	Y	Integer	Single	0 = None 1 = Policy defined by remote FW Are-You-There (AYT) 2 = Policy pushed CPP 4 = Policy from server
IPSec-Sec-Association	Y			String	Single	Name of the security association
IPSec-Split-DNS-Names	Y	Y	Y	String	Single	Specifies the list of secondary domain names to send to the client (1 - 255 characters).
IPSec-Split-Tunneling-Policy	Y	Y	Y	Integer	Single	0 = Tunnel everything 1 = Split tunneling 2 = Local LAN permitted
IPSec-Split-Tunnel-List	Y	Y	Y	String	Single	Specifies the name of the network or access list that describes the split tunnel inclusion list.
IPSec-Tunnel-Type	Y	Y	Y	Integer	Single	1 = LAN-to-LAN 2 = Remote access
IPSec-User-Group-Lock	Y			Boolean	Single	0 = Disabled 1 = Enabled

Table C-2 Security Appliance Supported Cisco Attributes for LDAP Authorization (continued)

Attribute Name/	VPN 3000	ASA	PIX	Syntax/ Type	Single or Multi-Valued	Possible Values
L2TP-Encryption	Y			Integer	Single	Bitmap: 1 = Encryption required 2 = 40 bit 4 = 128 bits 8 = Stateless-Req 15 = 40/128-Encr/Stateless-Req
L2TP-MPPC-Compression	Y			Integer	Single	0 = Disabled 1 = Enabled
MS-Client-Subnet-Mask	Y	Y	Y	String	Single	An IP address
PFS-Required	Y	Y	Y	Boolean	Single	0 = No 1 = Yes
Port-Forwarding-Name	Y	Y		String	Single	Name string (for example, “Corporate-Apps”)
PPTP-Encryption	Y			Integer	Single	Bitmap: 1 = Encryption required 2 = 40 bits 4 = 128 bits 8 = Stateless-Required Example: 15 = 40/128-Encr/Stateless-Req
PPTP-MPPC-Compression	Y			Integer	Single	0 = Disabled 1 = Enabled
Primary-DNS	Y	Y	Y	String	Single	An IP address
Primary-WINS	Y	Y	Y	String	Single	An IP address
Privilege-Level						
Required-Client-Firewall-Vendor-Code	Y	Y	Y	Integer	Single	1 = Cisco Systems (with Cisco Integrated Client) 2 = Zone Labs 3 = NetworkICE 4 = Sygate 5 = Cisco Systems (with Cisco Intrusion Prevention Security Agent)
Required-Client-Firewall-Description	Y	Y	Y	String	Single	String

Table C-2 Security Appliance Supported Cisco Attributes for LDAP Authorization (continued)

Attribute Name/	VPN 3000	ASA	PIX	Syntax/ Type	Single or Multi-Valued	Possible Values
Required-Client-Firewall-Product-Code	Y	Y	Y	Integer	Single	Cisco Systems Products: 1 = Cisco Intrusion Prevention Security Agent or Cisco Integrated Client (CIC) Zone Labs Products: 1 = Zone Alarm 2 = Zone AlarmPro 3 = Zone Labs Integrity NetworkICE Product: 1 = BlackIce Defender/Agent Sygate Products: 1 = Personal Firewall 2 = Personal Firewall Pro 3 = Security Agent
Require-HW-Client-Auth	Y	Y	Y	Boolean	Single	0 = Disabled 1 = Enabled
Require-Individual-User-Auth	Y	Y	Y	Integer	Single	0 = Disabled 1 = Enabled
Secondary-DNS	Y	Y	Y	String	Single	An IP address
Secondary-WINS	Y	Y	Y	String	Single	An IP address
SEP-Card-Assignment				Integer	Single	Not used
Simultaneous-Logins	Y	Y	Y	Integer	Single	0-2147483647
Strip-Realm	Y	Y	Y	Boolean	Single	0 = Disabled 1 = Enabled
TACACS-Authtype	Y	Y	Y	Interger	Single	
TACACS-Privilege-Level	Y	Y	Y	Interger	Single	
Tunnel-Group-Lock		Y	Y	String	Single	Name of the tunnel group or “none”
Tunneling-Protocols	Y	Y	Y	Integer	Single	1 = PPTP 2 = L2TP 4 = IPSec 8 = L2TP/IPSec 16 = WebVPN. 8 and 4 are mutually exclusive (0 - 11, 16 - 27 are legal values)
Use-Client-Address	Y			Boolean	Single	0 = Disabled 1 = Enabled
User-Auth-Server-Name	Y			String	Single	IP address or hostname
User-Auth-Server-Port	Y			Integer	Single	Port number for server protocol
User-Auth-Server-Secret	Y			String	Single	Server password
WebVPN-ACL-Filters		Y		String	Single	Webtype Access-List name

Table C-2 Security Appliance Supported Cisco Attributes for LDAP Authorization (continued)

Attribute Name/	VPN 3000	ASA	PIX	Syntax/ Type	Single or Multi-Valued	Possible Values
WebVPN-Apply-ACL-Enable	Y	Y		Integer	Single	0 = Disabled 1 = Enabled With version 8.0 and later, this attribute is not required.
WebVPN-Citrix-Support-Enable	Y	Y		Integer	Single	0 = Disabled 1 = Enabled With version 8.0 and later, this attribute is not required.
WebVPN-Enable-functions				Integer	Single	Not used - deprecated
WebVPN-Exchange-Server-Address				String	Single	Not used - deprecated
WebVPN-Exchange-Server-NETBIOS-Name				String	Single	Not used - deprecated
WebVPN-File-Access-Enable	Y	Y		Integer	Single	0 = Disabled 1 = Enabled
WebVPN-File-Server-Browsing-Enable	Y	Y		Integer	Single	0 = Disabled 1 = Enabled
WebVPN-File-Server-Entry-Enable	Y	Y		Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Forwarded-Ports		Y		String	Single	Port-Forward list name
WebVPN-Homepage	Y	Y		String	Single	A URL such as http://example-portal.com.
WebVPN-Macro-Substitution-Value1	Y	Y		String	Single	See *SSL VPN Deployment Guide* for examples and use cases at this URL: http://supportwiki.cisco.com/ViewWiki/index.php/Cisco_ASA_5500_SSL_VPN_Deployment_Guide%2C_Version_8.x
WebVPN-Macro-Substitution-Value2	Y	Y		String	Single	See *SSL VPN Deployment Guide* for examples and use cases at this URL: http://supportwiki.cisco.com/ViewWiki/index.php/Cisco_ASA_5500_SSL_VPN_Deployment_Guide%2C_Version_8.x
WebVPN-Port-Forwarding-Auto-Download-Enable	Y	Y		Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Port-Forwarding- Enable	Y	Y		Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Port-Forwarding-Exchange-Proxy-Enable	Y	Y		Integer	Single	0 = Disabled 1 = Enabled

Table C-2 Security Appliance Supported Cisco Attributes for LDAP Authorization (continued)

Attribute Name/	VPN 3000	ASA	PIX	Syntax/ Type	Single or Multi-Valued	Possible Values
WebVPN-Port-Forwarding-HTTP-Proxy-Enable	Y	Y		Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Single-Sign-On-Server-Name		Y		String	Single	Name of the SSO Server (1 - 31 characters).
WebVPN-SVC-Client-DPD	Y	Y		Integer	Single	0 = Disabled n = Dead Peer Detection value in seconds (30 - 3600)
WebVPN-SVC-Compression	Y	Y		Integer	Single	0 = None 1 = Deflate Compression
WebVPN-SVC-Enable	Y	Y		Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SVC-Gateway-DPD	Y	Y		Integer	Single	0 = Disabled n = Dead Peer Detection value in seconds (30 - 3600)
WebVPN-SVC-Keepalive	Y	Y		Integer	Single	0 = Disabled n = Keepalive value in seconds (15 - 600)
WebVPN-SVC-Keep-Enable	Y	Y		Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SVC-Rekey-Method	Y	Y		Integer	Single	0 = None 1 = SSL 2 = New tunnel 3 = Any (sets to SSL)
WebVPN-SVC-Rekey-Period	Y	Y		Integer	Single	0 = Disabled n = Retry period in minutes (4 - 10080)
WebVPN-SVC-Required-Enable	Y	Y		Integer	Single	0 = Disabled 1 = Enabled
WebVPN-URL-Entry-Enable	Y	Y		Integer	Single	0 = Disabled 1 = Enabled
WebVPN-URL-List		Y		String	Single	URL-list name

Cisco AV Pair Attribute Syntax

The Cisco Attribute Value (AV) pair (ID# 26/9/1) can be used to enforce access lists from a Radius server (like Cisco ACS), or from an LDAP server via an ldap-attribute-map.

The syntax of each Cisco-AV-Pair rule is as follows:

[Prefix] [Action] [Protocol] [Source] [Source Wildcard Mask] [Destination] [Destination Wildcard Mask] [Established] [Log] [Operator] [Port]

Table C-3 describes the syntax rules.

Table C-3 AV-Pair Attribute Syntax Rules

Field	Description
Prefix	A unique identifier for the AV pair. For example: `ip:inacl#1=` (for standard access lists) or `webvpn:inacl#` (for clientless SSL VPN access lists). This field only appears when the filter has been sent as an AV pair.
Action	Action to perform if rule matches: deny, permit.
Protocol	Number or name of an IP protocol. Either an integer in the range 0 - 255 or one of the following keywords: icmp, igmp, ip, tcp, udp.
Source	Network or host that sends the packet. Specify it as an IP address, a hostname, or the keyword "any." If using an IP address, the source wildcard mask must follow. This field does not apply to Clientless SSL VPN because the security appliance plays the role of the source/proxy
Source Wildcard Mask	The wildcard mask that applies to the source address. This field does not apply to Clientless SSL VPN because the security appliance plays the role of the source/proxy
Destination	Network or host that receives the packet. Specify as an IP address, a hostname, or the keyword "any." If using an IP address, the source wildcard mask must follow.
Destination Wildcard Mask	The wildcard mask that applies to the destination address.
Log	Generates a FILTER log message. You must use this keyword to generate events of severity level 9.
Operator	Logic operators: greater than, less than, equal to, not equal to.
Port	The number of a TCP or UDP port in the range 0 - 65535.

Cisco AV Pairs ACL Examples

Table C-4 shows examples of Cisco AV pairs and describes the allow or deny actions that result.

Note Each ACL # in `inacl#` must be unique. However, they do not need to be sequential (i.e. 1, 2, 3, 4). For example, they could be 5, 45, 135.

Table C-4 Examples of Cisco AV Pairs and their Permitting or Denying Action

Cisco AV Pair Example	Permitting or Denying Action
`ip:inacl#1=deny ip 10.155.10.0 0.0.0.255 10.159.2.0 0.0.0.255 log`	Allows IP traffic between the two hosts using full tunnel IPsec or SSL VPN client.
`ip:inacl#2=permit TCP any host 10.160.0.1 eq 80 log`	Allows TCP traffic from all hosts to the specific host on port 80 only using full tunnel IPsec or SSL VPN client.

`webvpn:inacl#1=permit url http://www.website.com`
`webvpn:inacl#2=deny url smtp://server`
`webvpn:inacl#3=permit url cifs://server/share`
	Allows clientless traffic to the URL specified, denies smtp traffic to a specific server, and allows file share access (CIFS) to the specified server.

`webvpn:inacl#1=permit tcp 10.86.1.2 eq 2222 log`
`webvpn:inacl#2=deny tcp 10.86.1.2 eq 2323 log`
	Denies telnet and permits SSH on non-default ports 2323 and 2222, respectively.

`webvpn:inacl#1=permit url ssh://10.86.1.2`
`webvpn:inacl#35=permit tcp 10.86.1.5 eq 22 log`
`webvpn:inacl#48=deny url telnet://10.86.1.2`
`webvpn:inacl#100=deny tcp 10.86.1.6 eq 23`
	Allows SSH to default port 22 and 23, respectively. For this example, we assume you are using telnet/ssh java plugins enforced by these ACLs.

URL Types supported in ACLs

The URL may be a partial URL, contain wildcards for the server, or contain a port.

The following URL types are supported:

any All URLs	http://	nfs://	sametime://	telnet://
cifs://	https://	pop3://	smart-tunnel://	tn3270://
citrix://	ica://	post://	smtp://	tn5250://
citrixs://	imap4://	rdp://	ssh://	vnc://
ftp://				

Note The URLs listed above appear in CLI or ASDM menus based on whether the associated plugin is enabled.

Guidelines for using Cisco-AV Pairs (ACLs)

- Use Cisco-AV pair entries with the ip:inacl# prefix to enforce access lists for remote IPSec and SSL VPN Client (SVC) tunnels.
- Use Cisco-AV pair entries with the webvpn:inacl# prefix to enforce access lists for SSL VPN clientless (browser-mode) tunnels.
- For Webtype ACLs, you don’t specify the source because the security appliance is the source.

Table C-5 lists the tokens for the Cisco-AV-pair attribute:

Table C-5 Security Appliance-Supported Tokens

Token	Syntax Field	Description
ip:inacl#*Num*=	N/A (Identifier)	(Where *Num* is a unique integer.) Starts all AV pair access control lists. Enforces access lists for remote IPSec and SSL VPN (SVC) tunnels.
webvpn:inacl#*Num*=	N/A (Identifier)	(Where *Num* is a unique integer.) Starts all clientless SSL AV pair access control lists. Enforces access lists for clientless (browser-mode) tunnels.
deny	Action	Denies action. (Default)
permit	Action	Allows action.
icmp	Protocol	Internet Control Message Protocol (ICMP)
1	Protocol	Internet Control Message Protocol (ICMP)
IP	Protocol	Internet Protocol (IP)
0	Protocol	Internet Protocol (IP)
TCP	Protocol	Transmission Control Protocol (TCP)
6	Protocol	Transmission Control Protocol (TCP)
UDP	Protocol	User Datagram Protocol (UDP)
17	Protocol	User Datagram Protocol (UDP)
any	Hostname	Rule applies to any host.
host	Hostname	Any alpha-numeric string that denotes a hostname.
log	Log	When the event is hit, a filter log message appears. (Same as permit and log or deny and log.)
lt	Operator	Less than value
gt	Operator	Greater than value
eq	Operator	Equal to value
neq	Operator	Not equal to value
range	Operator	Inclusive range. Should be followed by two values.

Active Directory/LDAP VPN Remote Access Authorization Use Cases

This section presents example procedures for configuring authentication and authorization on the security appliance using the Microsoft Active Directory server. It includes the following use cases:

- User-Based Attributes Policy Enforcement, page C-18
- Placing LDAP users in a specific Group-Policy, page C-20
- Enforcing Static IP Address Assignment for AnyConnect Tunnels, page C-22
- Enforcing Dial-in Allow or Deny Access, page C-25
- Enforcing Logon Hours and Time-of-Day Rules, page C-28

Other configuration examples available on Cisco.com include the following TechNotes:

- *ASA/PIX: Mapping VPN Clients to VPN Group Policies Through LDAP Configuration Example* at:

http://www.cisco.com/en/US/products/ps6120/products_configuration_example09186a008089149d.shtml

- *PIX/ASA 8.0: Use LDAP Authentication to Assign a Group Policy at Login* at:

  http://www.cisco.com/en/US/partner/products/ps6120/products_configuration_example09186a00808d1a7c.shtml

User-Based Attributes Policy Enforcement

Any standard LDAP attribute can be mapped to a well-known Vendor Specific Attribute (VSA) Likewise, one or more LDAP attribute(s) can be mapped to one or more Cisco LDAP attributes.

In this use case we configure the security appliance to enforce a simple banner for a user configured on an AD LDAP server. For this case, on the server, we use the Office field in the General tab to enter the banner text. This field uses the attribute named *physicalDeliveryOfficeName*. On the security appliance, we create an attribute map that maps *physicalDeliveryOfficeName* to the Cisco attribute *Banner1*. During authentication, the security appliance retrieves the value of physicalDeliveryOfficeName from the server, maps the value to the Cisco attribute Banner1, and displays the banner to the user.

This case applies to any connection type, including the IPSec VPN client, AnyConnect SSL VPN client, or clientless SSL VPN. For the purposes of this case, User1 is connecting through a clientless SSL VPN connection.

Step 1 Configure the attributes for a user on the AD/LDAP Server.

Right-click a user. The properties window displays (Figure C-3). Click the General tab and enter some banner text in the Office field. The Office field uses the AD/LDAP attribute *physicalDeliveryOfficeName*.

Figure C-3 Figure 3 LDAP User configuration

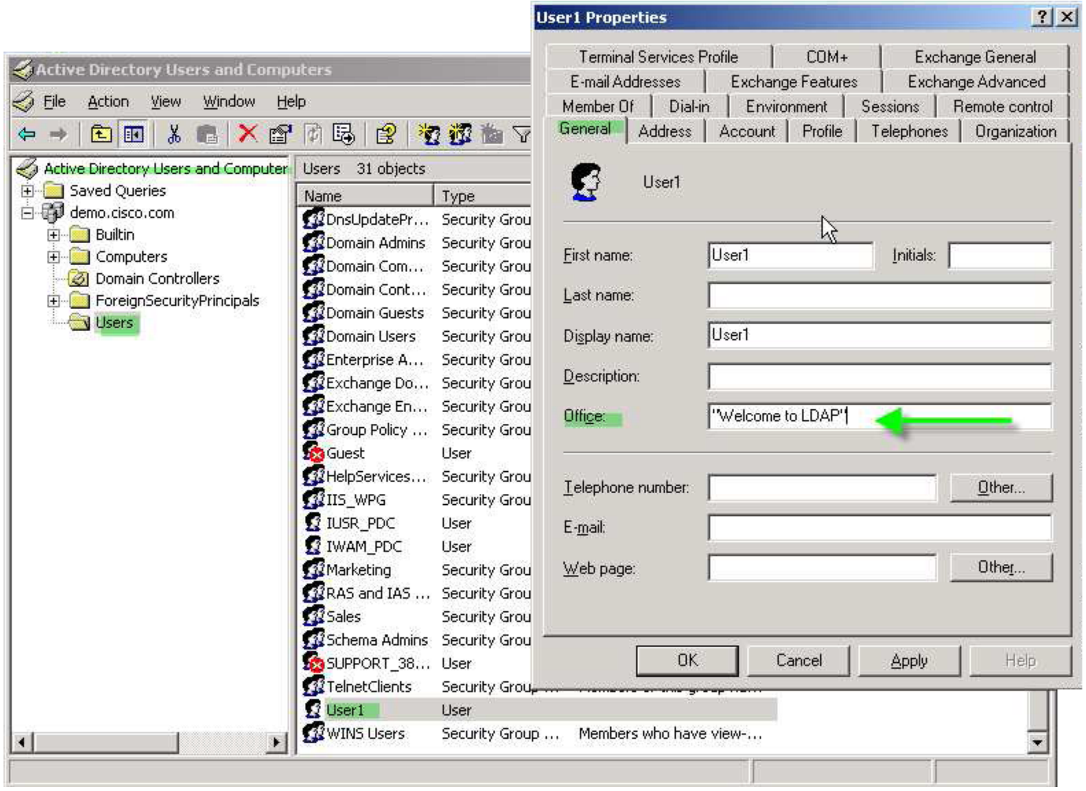

Name	Type
DnsUpdatePr…	Security Grou
Domain Admins	Security Grou
Domain Com…	Security Grou
Domain Cont…	Security Grou
Domain Guests	Security Grou
Domain Users	Security Grou
Enterprise A…	Security Grou
Exchange Do…	Security Grou
Exchange En…	Security Grou
Group Policy …	Security Grou
Guest	User
HelpServices…	Security Grou
IIS_WPG	Security Grou
IUSR_PDC	User
IWAM_PDC	User
Marketing	Security Grou
RAS and IAS …	Security Grou
Sales	Security Grou
Schema Admins	Security Grou
SUPPORT_38…	User
TelnetClients	Security Group …
User1	User
WINS Users	Security Group … Members who have view-…

First name: User1 Initials:
Last name:
Display name: User1
Description:
Office: "Welcome to LDAP"
Telephone number:
E-mail:
Web page:

Step 2 Create an LDAP attribute map on the security appliance:

The following example creates the map *Banner*, and maps the AD/LDAP attribute *physicalDeliveryOfficeName* to the Cisco attribute *Banner1*:

```
hostname(config)# ldap attribute-map Banner
hostname(config-ldap-attribute-map)# map-name physicalDeliveryOfficeName Banner1
```

Step 3 Associate the LDAP attribute map to the AAA server.

The following example enters the aaa server host configuration more for the host *3.3.3.4*, in the AAA server group *MS_LDAP*, and associates the attribute map *Banner* that you created in step 2:

```
hostname(config)# aaa-server MS_LDAP host 3.3.3.4
hostname(config-aaa-server-host)# ldap-attribute-map Banner
```

Step 4 Test the banner enforcement.

This example shows a clientless SSL connection and the banner enforced through the attribute map after the user authenticates (Figure C-4).

Figure C-4 Banner Displayed

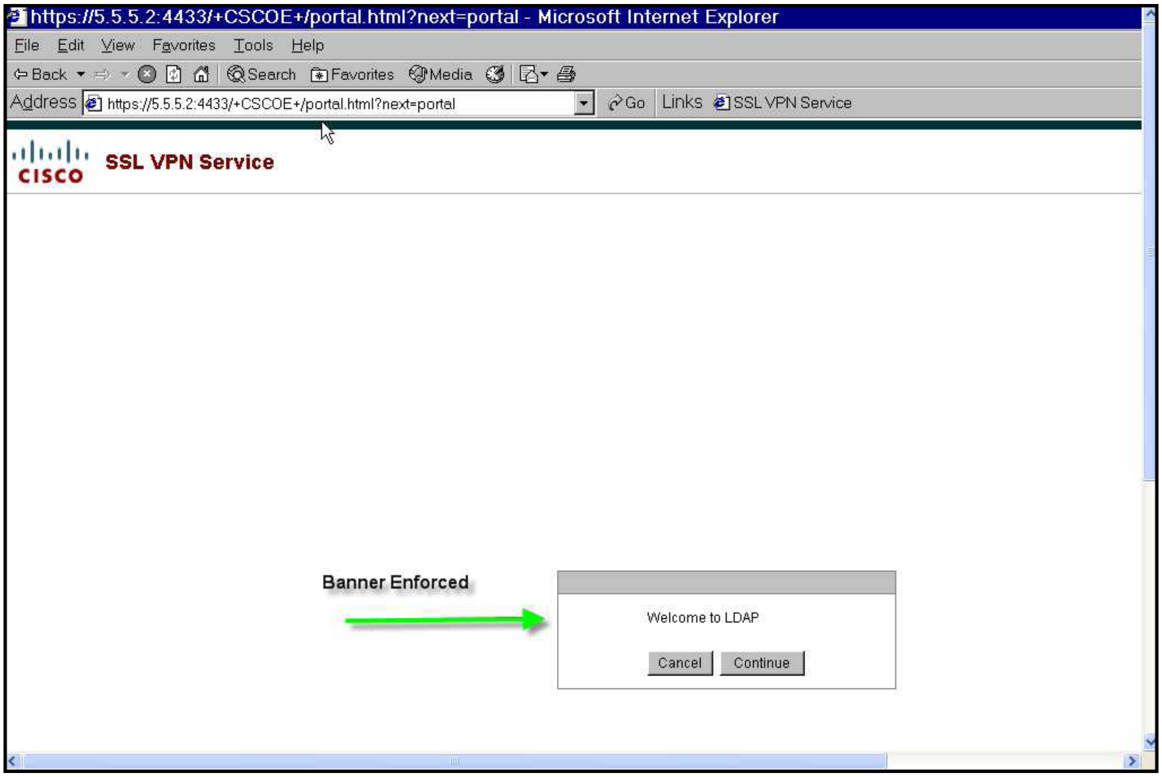
https://5.5.5.2:4433/+CSCOE+/portal.html?next=portal - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address https://5.5.5.2:4433/+CSCOE+/portal.html?next=portal
Go Links SSL VPN Service
CISCO
SSL VPN Service
Banner Enforced
Welcome to LDAP
Cancel
Continue

Placing LDAP users in a specific Group-Policy

In this case we authenticate User1 on the AD LDAP server to a specific group policy on the security appliance. On the server, we use the *Department* field of the Organization tab to enter the name of the group policy. Then we create an attribute map and map Department to the Cisco attribute *IETF-Radius-Class*. During authentication, the security appliance retrieves the value of Department from the server, maps the value to the IETF-Radius-Class, and places User1 in the group policy.

This case applies to any connection type, including the IPSec VPN client, AnyConnect SSL VPN client, or clientless SSL VPN. For the purposes of this case, user1 is connecting through a clientless SSL VPN connection.

Step 1 Configure the attributes for the user on the AD LDAP Server.

Right-click the user. The Properties window displays (Figure C-5). Click the Organization tab and enter *Group-Policy-1* in the Department field.

Figure C-5 AD LDAP Department attribute

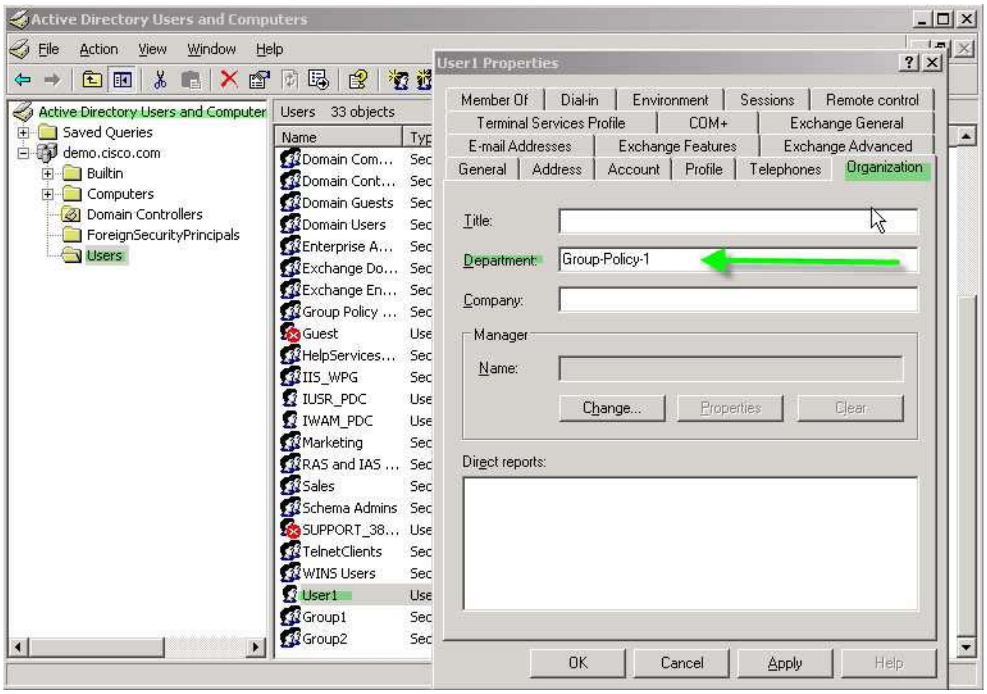

Department: Group-Policy-1

Step 2 Define an attribute map for the LDAP configuration shown in Step 1.

In this case we map the AD attribute Department to the Cisco attribute IETF-Radius-Class. For example:

```
hostname(config)# ldap attribute-map group_policy
hostname(config-ldap-attribute-map)# map-name Department IETF-Radius-Class
```

Step 3 Associate the LDAP attribute map to the AAA server.

The following example enters the aaa server host configuration mode for the host *3.3.3.4*, in the AAA server group *MS_LDAP*, and associates the attribute map *group_policy* that you created in step 2:

```
hostname(config)# aaa-server MS_LDAP host 3.3.3.4
hostname(config-aaa-server-host)# ldap-attribute-map group_policy
```

Step 4 Add the new group-policy on the security appliance and configure the required policy attributes that will be assigned to the user. For this case, we created the Group-policy-1, the name entered in the Department field on the server:

```
hostname(config)# group-policy Group-policy-1 external server-group LDAP_demo
hostname(config-aaa-server-group)#
```

Step 5 Establish the VPN connection as the user would, and verify that the session inherits the attributes from Group-Policy1 (and any other applicable attributes from the default group-policy)

You can monitor the communication between the security appliance and the server by enabling the **debug ldap 255** command from privileged EXEC mode. Below is sample output of this command. The output has been edited to provide the key messages:

[29] Authentication successful for user1 to 3.3.3.4

[29] Retrieving user attributes from server 3.3.3.4

[29] Retrieved Attributes:

[29] department: value = Group-Policy-1

[29] mapped to IETF-Radius-Class: value = Group-Policy-1

Enforcing Static IP Address Assignment for AnyConnect Tunnels

In this case we configure the AnyConnect client user *Web1* to receive a static IP Address. We enter the address in the *Assign Static IP Address* field of the Dialin tab on the AD LDAP server. This field uses the *msRADIUSFramedIPAddress* attribute. We create an attribute map that maps it to the Cisco attribute *IETF-Radius-Framed-IP-Address*.

During authentication, the security appliance retrieves the value of msRADIUSFramedIPAddress from the server, maps the value to the Cisco attribute IETF-Radius-Framed-IP-Address, and provides the static address to User1 .

This case applies to full-tunnel clients, including the IPSec client and the SSL VPN clients (AnyConnect client 2.x and the legacy SSL VPN client).

Step 1 Configure the user attributes on the AD LDAP server.

Right-click on the user name. The Properties window displays (Figure C-6). Click the Dialin tab, check *Assign Static IP Address*, and enter an IP address. For this case we use 3.3.3.233.

Figure C-6 Assign Static IP Address

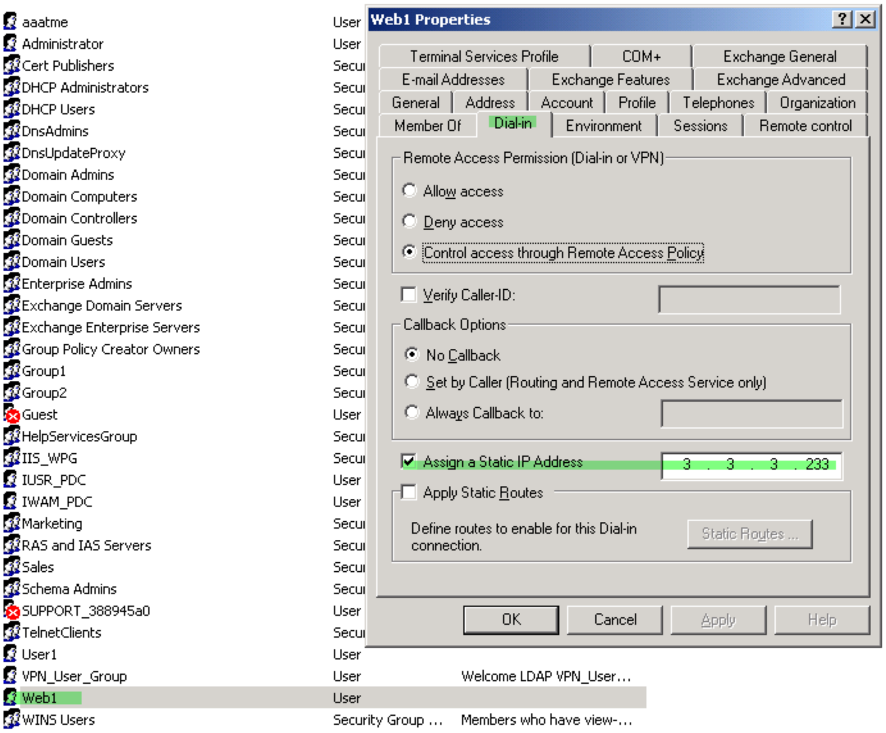

aaatme	User	
Administrator	User	
Cert Publishers	Secu	
DHCP Administrators	Secu	
DHCP Users	Secu	
DnsAdmins	Secu	
DnsUpdateProxy	Secu	
Domain Admins	Secu	
Domain Computers	Secu	
Domain Controllers	Secu	
Domain Guests	Secu	
Domain Users	Secu	
Enterprise Admins	Secu	
Exchange Domain Servers	Secu	
Exchange Enterprise Servers	Secu	
Group Policy Creator Owners	Secu	
Group1	Secu	
Group2	Secu	
Guest	User	
HelpServicesGroup	Secu	
IIS_WPG	Secu	
IUSR_PDC	User	
IWAM_PDC	User	
Marketing	Secu	
RAS and IAS Servers	Secu	
Sales	Secu	
Schema Admins	Secu	
SUPPORT_388945a0	User	
TelnetClients	Secu	
User1	User	
VPN_User_Group	User	Welcome LDAP VPN_User...
Web1	User	
WINS Users	Security Group ...	Members who have view-...

Web1 Properties

Terminal Services Profile | COM+ | Exchange General
E-mail Addresses | Exchange Features | Exchange Advanced
General | Address | Account | Profile | Telephones | Organization
Member Of | Dial-in | Environment | Sessions | Remote control

Remote Access Permission (Dial-in or VPN)

○ Allow access
○ Deny access
● Control access through Remote Access Policy

□ Verify Caller-ID:

Callback Options

● No Callback
○ Set by Caller (Routing and Remote Access Service only)
○ Always Callback to:

☑ Assign a Static IP Address 3 . 3 . 3 . 233

□ Apply Static Routes

Define routes to enable for this Dial-in connection.

Step 2 Create an attribute map for the LDAP configuration shown in Step 1.

In this case we map the AD attribute *msRADIUSFrameIPAddress* used by the Static Address field to the Cisco attribute *IETF-Radius-Framed-IP-Address*.

For example:

```
hostname(config)# ldap attribute-map static_address
hostname(config-ldap-attribute-map)# map-name msRADIUSFrameIPAddress
IETF-Radius-Framed-IP-Address
```

Step 3 Associate the LDAP attribute map to the AAA server.

The following example enters the aaa server host configuration mode for the host *3.3.3.4*, in the AAA server group *MS_LDAP*, and associates the attribute map *static_address* that you created in step 2:

```
hostname(config)# aaa-server MS_LDAP host 3.3.3.4
hostname(config-aaa-server-host)# ldap-attribute-map static_address
```

Step 4 Verify the **vpn-address-assigment** command is configured to specify aaa by viewing this part of the configuration with the **show run all vpn-addr-assign command**:

vpn-addr-assign aaa

```
hostname(config)# show run all vpn-addr-assign
vpn-addr-assign aaa   <<<< ensure this configured.
no vpn-addr-assign dhcp
vpn-addr-assign local
hostname(config)#
```

Step 5 Establish a connection to the security appliance with the AnyConnect client. Observe the following:

- The banner is received in the same sequence as a clientless connection (Figure C-7).
- The user receives the IP address configured on the server and mapped to the security appliance (Figure C-8).

Figure C-7 Verify the Banner for the AnyConnect Session

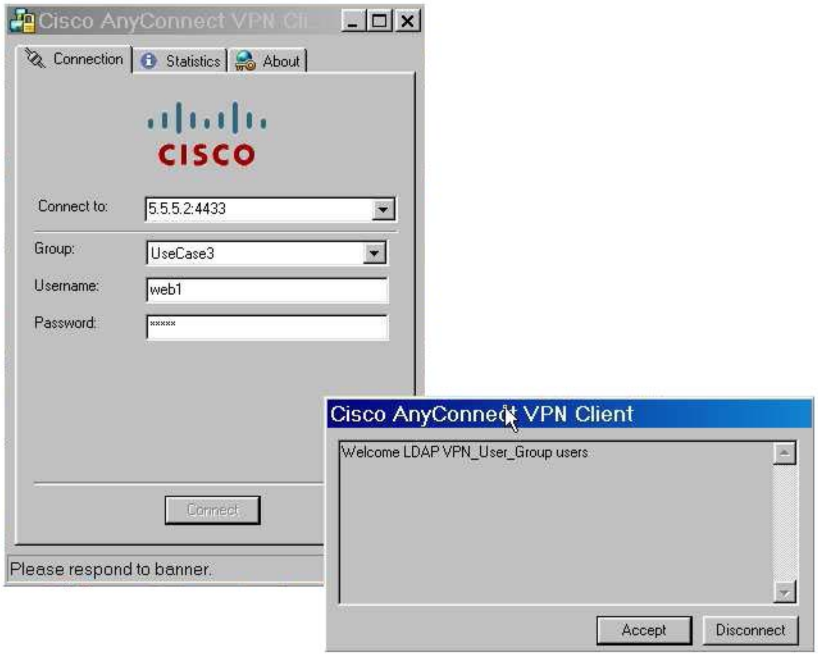

Connect to: 5.5.5.2:4433
Group: UseCase3
Username: web1
Password: *****
Please respond to banner.

Welcome LDAP VPN_User_Group users

Figure C-8 AnyConnect Session Established

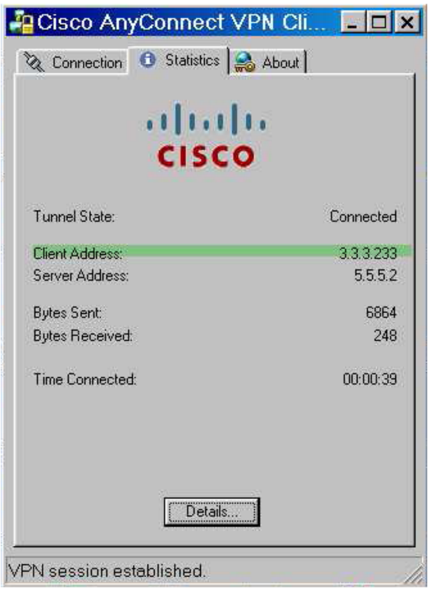

Tunnel State: Connected

Client Address: 3.3.3.233
Server Address: 5.5.5.2

Bytes Sent: 6864
Bytes Received: 248

Time Connected: 00:00:39

VPN session established.

You can use the **show vpn-sessiondb svc** command to view the session details and verify the address assigned:

```
hostname# show vpn-sessiondb svc

Session Type: SVC
Username     : web1                   Index        : 31
Assigned IP  : 3.3.3.233              Public IP    : 10.86.181.70
Protocol     : Clientless SSL-Tunnel DTLS-Tunnel
Encryption   : RC4 AES128             Hashing      : SHA1
Bytes Tx     : 304140                 Bytes Rx     : 470506
Group Policy : VPN_User_Group         Tunnel Group : UseCase3_TunnelGroup
Login Time   : 11:13:05 UTC Tue Aug 28 2007
Duration     : 0h:01m:48s
NAC Result   : Unknown
VLAN Mapping : N/A                    VLAN         : none

BXB-ASA5540#
```

Enforcing Dial-in Allow or Deny Access

In this case, we create an LDAP attribute map that specifies the tunneling protocols allowed by the user. We map the Allow Access and Deny Access settings on the Dialin tab to the Cisco attribute Tunneling-Protocols. The Cisco Tunneling-Protocols supports the bit-map values shown in Table C-6:

Table C-6 Bitmap Values for Cisco Tunneling-Protocol Attribute

Value	Tunneling Protocol
1	PPTP
2	L2TP
4[1]	IPSec
8[2]	L2TP/IPSEC
16	clientless SSL
32	SSL Client—AnyConnect or legacy SSL VPN client

[1] IPSec and L2TP over IPSec are not supported simultaneously. Therefore, the values 4 and 8 are mutually exclusive.

[2] See note 1.

Using this attribute, we create an Allow Access (TRUE) or a Deny Access (FALSE) condition for the protocols and enforce what method the user is allowed access with.

For this simplified example, by mapping the tunnel-protocol IPSec (4), we can create an allow (true) condition for the IPSec Client. We also map WebVPN (16) and SVC/AC (32) which is mapped as value of 48 (16+32) and create a deny (false) condition. This allows the user to connect to the security appliance using IPSec, but any attempt to connect using clientless SSL or the AnyConnect client is denied.

Another example of enforcing Dial-in Allow Acess or Deny Access can be found in the Tech Note *ASA/PIX: Mapping VPN Clients to VPN Group Policies Through LDAP Configuration Example*, at this URL:

> http://www.cisco.com/en/US/products/ps6120/products_configuration_example09186a008089149d.shtml

Step 1 Configure the user attributes on the AD LDAP server.

Right-click on the user. The Properties window displays. Click the Dial-in tab. Select **Allow Access** (Figure C-9).

Figure C-9 AD-LDAP user1 - Allow access

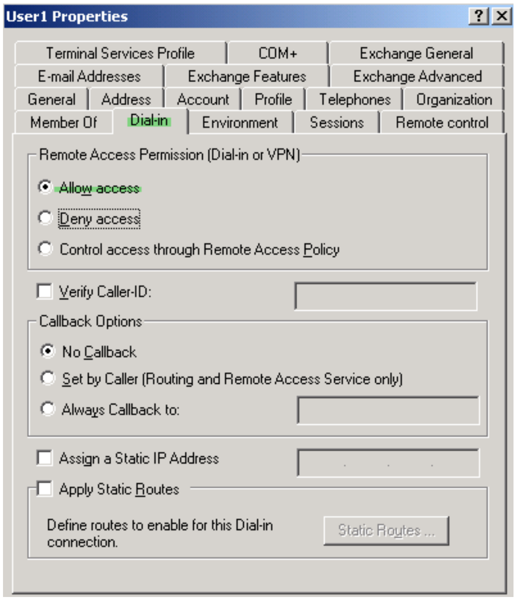

Remote Access Permission (Dial-in or VPN)

● Allow access
○ Deny access
○ Control access through Remote Access Policy

□ Verify Caller-ID:

Callback Options

● No Callback
○ Set by Caller (Routing and Remote Access Service only)
○ Always Callback to:

□ Assign a Static IP Address

□ Apply Static Routes

Define routes to enable for this Dial-in connection.

Note If you select the third option "Control access through the Remote Access Policy", then a value is not returned from the server, and the permissions that are enforced are based on the internal group policy settings of the security appliance.

Step 2 Create an attribute map to allow both an IPSec and AnyConnect connection, but deny a clientless SSL connection.

In this case we create the map *tunneling_protocols*, and map the AD attribute *msNPAllowDialin* used by the Allow Access setting to the Cisco attribute *Tunneling-Protocols* using the **map-name** command, and add map values with the **map-value** command,

For example:

```
hostname(config)# ldap attribute-map tunneling_protocols
hostname(config-ldap-attribute-map)# map-name msNPAllowDialin Tunneling-Protocols
hostname(config-ldap-attribute-map)# map-value msNPAllowDialin FALSE  48
hostname(config-ldap-attribute-map)# map-value msNPAllowDialin TRUE 4
```

Step 3 Associate the LDAP attribute map to the AAA server.

The following example enters the aaa server host configuration mode for the host *3.3.3.4*, in the AAA server group *MS_LDAP*, and associates the attribute map *tunneling_protocols* that you created in step 2:

```
hostname(config)# aaa-server MS_LDAP host 3.3.3.4
hostname(config-aaa-server-host)# ldap-attribute-map tunneling_protocols
```

Step 4 Verify the attribute map works as configured.

Using a PC as a remote user would, attempt connections using clientless SSL, the AnyConnect client, and the IPSec client. The clientless and AnyConnect connections should fail and the user should be informed that an unauthorized connection mechanism was the reason for the failed connection. The IPSec client should connect because IPSec is an allowed tunneling protocol according to attribute map.

Figure C-10 Login Denied Message for Clientless User

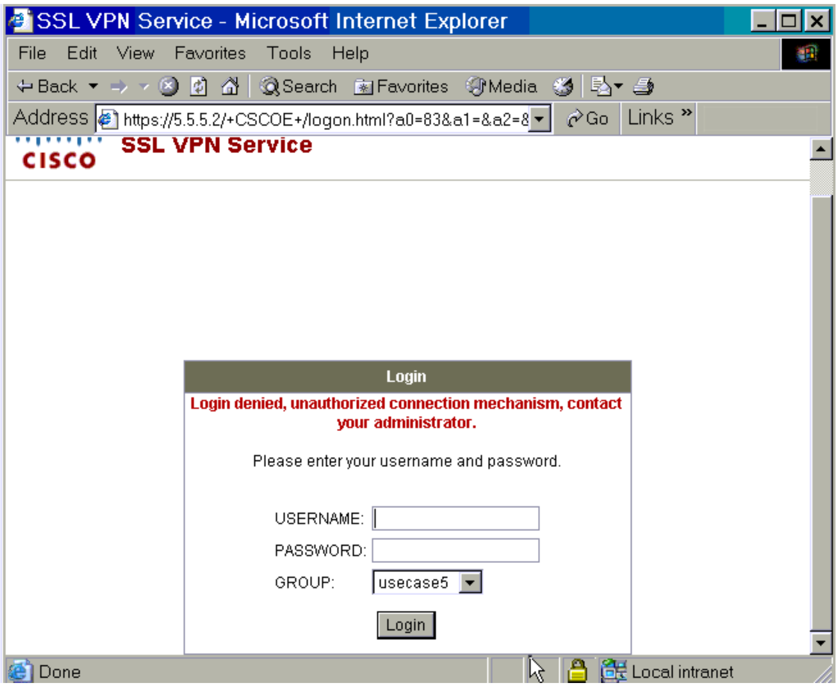

SSL VPN Service - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address https://5.5.5.2/+CSCOE+/logon.html?a0=83&a1=&a2=8
Go Links
CISCO SSL VPN Service
Login

Login denied, unauthorized connection mechanism, contact your administrator.

Please enter your username and password.

USERNAME:
PASSWORD:
GROUP: usecase5

Login
Done
Local intranet

Figure C-11 Login Denied Message for AnyConnect Client User.

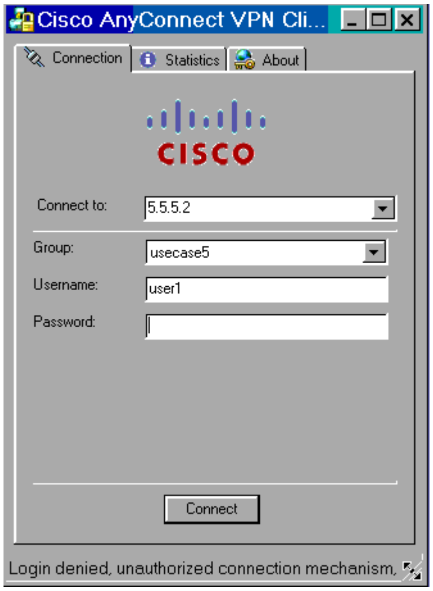

Connect to: 5.5.5.2
Group: usecase5
Username: user1
Password:

Login denied, unauthorized connection mechanism,

Enforcing Logon Hours and Time-of-Day Rules

In this use case we configure and enforce the hours that a clientless SSL user is allowed to access the network. A good example of this is when you want to allow a business partner access to the network only during normal business hours.

For this case, on the AD server, we use the *Office* field to enter the name of the partner. This field uses the *physicalDeliveryOfficeName* attribute. Then we create an attribute map on the security appliance to map that attribute to the Cisco attribute *Access-Hours*. During authentication, the security appliance retrieves the value of physicalDeliveryOfficeName (the Office field*)* and maps it to Access-Hours.

Step 1 Configure the user attributes on the AD LDAP server.

Select the user. Right click on Properties. The Properties window displays (Figure C-12). For this case, we use the Office field of the General tab:

Figure C-12 Active Directory - Time-range

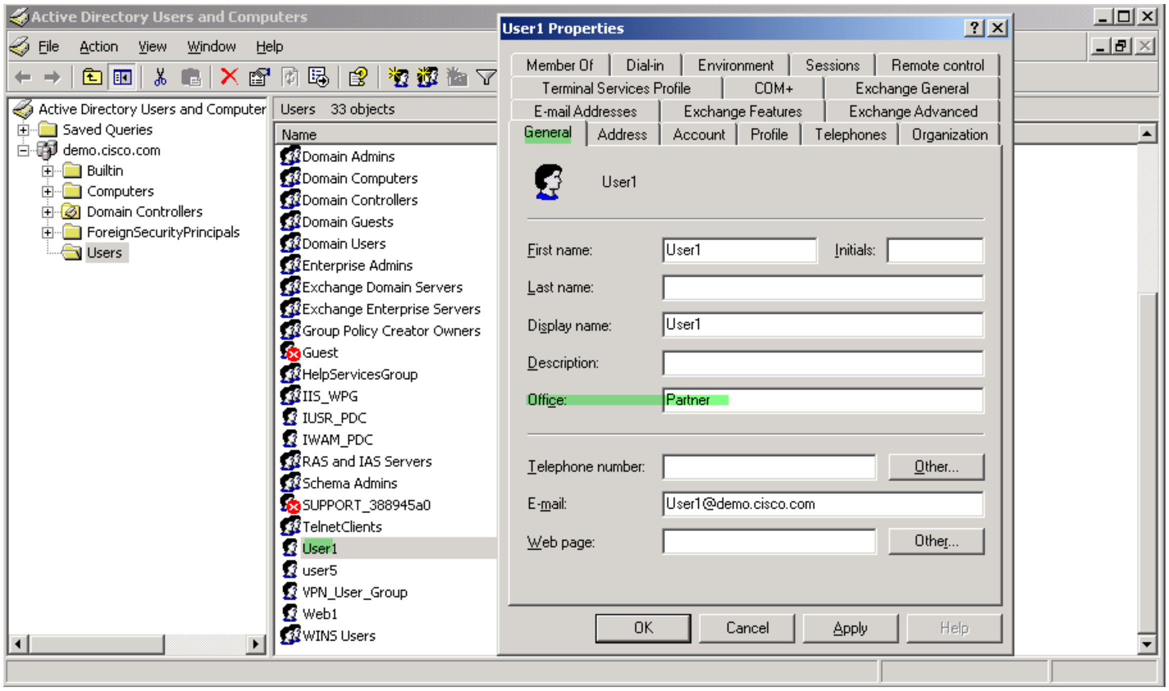

User1

First name: User1 Initials:

Last name:

Display name: User1

Description:

Office: Partner

Telephone number:

E-mail: User1@demo.cisco.com

Web page:

Step 2 Create an attribute map.

In this case we create the attribute map access_hours and map the AD attribute *physicalDeliveryOfficeName* used by the Office field to the Cisco attribute *Access-Hours*.

For example:

```
hostname(config)# ldap attribute-map access_hours
hostname(config-ldap-attribute-map)# map-name physicalDeliveryOfficeName Access-Hours
```

Step 3 Associate the LDAP attribute map to the AAA server.

The following example enters the aaa server host configuration mode for the host *3.3.3.4*, in the AAA server group *MS_LDAP*, and associates the attribute map *access_hours* that you created in step 2:

```
hostname(config)# aaa-server MS_LDAP host 3.3.3.4
hostname(config-aaa-server-host)# ldap-attribute-map access_hours
```

Step 4 Configure time ranges for each value allowed on the server. In this case, we entered Partner in the Office field for User1. Therefore, there must be a time range configured for Partner. The following example configures Partner access hours from 9am to 5pm Monday through Friday:

```
hostname(config)# time-range Partner
hostname(config-time-range)# periodic weekdays 09:00 to 17:00
```

Configuring an External RADIUS Server

This section presents an overview of the RADIUS configuration procedure and defines the Cisco RADIUS attributes. It includes the following topics:

- Reviewing the RADIUS Configuration Procedure, page C-30
- Security Appliance RADIUS Authorization Attributes, page C-30
- Security Appliance IETF RADIUS Authorization Attributes, page C-38

Reviewing the RADIUS Configuration Procedure

This section describes the RADIUS configuration steps required to support authentication and authorization of the security appliance users. Follow these steps to set up the RADIUS server to inter operate with the security appliance.

Step 1 Load the security appliance attributes into the RADIUS server. The method you use to load the attributes depends on which type of RADIUS server you are using:

- If you are using Cisco ACS: the server already has these attributes integrated. You can skip this step.
- If you are using a FUNK RADIUS server: Cisco supplies a dictionary file that contains all the security appliance attributes. Obtain this dictionary file, `cisco3k.dct`, from Software Center on CCO or from the security appliance CD-ROM. Load the dictionary file on your server.
- For other vendors' RADIUS servers (for example, Microsoft Internet Authentication Service): you must manually define each security appliance attribute. To define an attribute, use the attribute name or number, type, value, and vendor code (3076). For a list of security appliance RADIUS authorization attributes and values, see

Step 2 Set up the users or groups with the permissions and attributes to send during IPSec or SSL tunnel establishment.

Security Appliance RADIUS Authorization Attributes

Authorization refers to the process of enforcing permissions or attributes. A RADIUS server defined as an authentication server enforces permissions or attributes if they are configured.

Table C-7 lists all the possible security appliance supported RADIUS attributes that can be used for user authorization.

Note RADIUS attribute names do not contain the cVPN3000 prefix. Cisco Secure ACS 4.x supports this new nomenclature, but attribute names in pre-4.0 ACS releases still include the cVPN3000 prefix. The appliances enforce the RADIUS attributes based on attribute numeric ID, not attribute name. LDAP attributes are enforced by their name, not by the ID.

Table C-7 Security Appliance Supported RADIUS Attributes and Values

Attribute Name	VPN 3000	ASA	PIX	Attr. #	Syntax/ Type	Single or Multi-Valued	Description or Value
Access-Hours	Y	Y	Y	1	String	Single	Name of the time range, for example, Business-hours
Simultaneous-Logins	Y	Y	Y	2	Integer	Single	An integer 0 to 2147483647
Primary-DNS	Y	Y	Y	5	String	Single	An IP address
Secondary-DNS	Y	Y	Y	6	String	Single	An IP address
Primary-WINS	Y	Y	Y	7	String	Single	An IP address
Secondary-WINS	Y	Y	Y	8	String	Single	An IP address
SEP-Card-Assignment				9	Integer	Single	Not used
Tunneling-Protocols	Y	Y	Y	11	Integer	Single	1 = PPTP 2 = L2TP 4 = IPSec 8 = L2TP/IPSec 16 = WebVPN 4 and 8 are mutually exclusive; 0-11 and 16-27 are legal values.
IPSec-Sec-Association	Y			12	String	Single	Name of the security association
IPSec-Authentication	Y			13	Integer	Single	0 = None 1 = RADIUS 2 = LDAP (authorization only) 3 = NT Domain 4 = SDI 5 = Internal 6 = RADIUS with Expiry 7 = Kerberos/Active Directory
Banner1	Y	Y	Y	15	String	Single	Banner string
IPSec-Allow-Passwd-Store	Y	Y	Y	16	Boolean	Single	0 = Disabled 1 = Enabled
Use-Client-Address	Y			17	Boolean	Single	0 = Disabled 1 = Enabled
PPTP-Encryption	Y			20	Integer	Single	Bitmap: 1 = Encryption required 2 = 40 bits 4 = 128 bits 8 = Stateless-Required 15= 40/128-Encr/Stateless-Req

Table C-7 Security Appliance Supported RADIUS Attributes and Values (continued)

Attribute Name	VPN 3000	ASA	PIX	Attr. #	Syntax/ Type	Single or Multi-Valued	Description or Value
L2TP-Encryption	Y			21	Integer	Single	Bitmap: 1 = Encryption required 2 = 40 bit 4 = 128 bits 8 = Stateless-Req 15= 40/128-Encr/Stateless-Req
Group-Policy		Y	Y	25	String	Single	Sets the group policy for the remote access VPN session. For version 8.2 and later, use this attribute instead of IETF-Radius-Class. You can use one of the three following formats: • <group policy name> • OU=<group policy name> • OU=<group policy name>;
IPSec-Split-Tunnel-List	Y	Y	Y	27	String	Single	Specifies the name of the network/access list that describes the split tunnel inclusion list
IPSec-Default-Domain	Y	Y	Y	28	String	Single	Specifies the single default domain name to send to the client (1-255 characters)
IPSec-Split-DNS-Names	Y	Y	Y	29	String	Single	Specifies the list of secondary domain names to send to the client (1-255 characters)
IPSec-Tunnel-Type	Y	Y	Y	30	Integer	Single	1 = LAN-to-LAN 2 = Remote access
IPSec-Mode-Config	Y	Y	Y	31	Boolean	Single	0 = Disabled 1 = Enabled
IPSec-User-Group-Lock	Y			33	Boolean	Single	0 = Disabled 1 = Enabled
IPSec-Over-UDP	Y	Y	Y	34	Boolean	Single	0 = Disabled 1 = Enabled
IPSec-Over-UDP-Port	Y	Y	Y	35	Integer	Single	4001 - 49151, default = 10000
Banner2	Y	Y	Y	36	String	Single	A banner string that is concatenated to the Banner1 string, if configured.
PPTP-MPPC-Compression	Y			37	Integer	Single	0 = Disabled 1 = Enabled

Table C-7 Security Appliance Supported RADIUS Attributes and Values (continued)

Attribute Name	VPN 3000	ASA	PIX	Attr. #	Syntax/ Type	Single or Multi-Valued	Description or Value
L2TP-MPPC-Compression	Y			38	Integer	Single	0 = Disabled 1 = Enabled
IPSec-IP-Compression	Y	Y	Y	39	Integer	Single	0 = Disabled 1 = Enabled
IPSec-IKE-Peer-ID-Check	Y	Y	Y	40	Integer	Single	1 = Required 2 = If supported by peer certificate 3 = Do not check
IKE-Keep-Alives	Y	Y	Y	41	Boolean	Single	0 = Disabled 1 = Enabled
IPSec-Auth-On-Rekey	Y	Y	Y	42	Boolean	Single	0 = Disabled 1 = Enabled
Required-Client- Firewall-Vendor-Code	Y	Y	Y	45	Integer	Single	1 = Cisco Systems (with Cisco Integrated Client) 2 = Zone Labs 3 = NetworkICE 4 = Sygate 5 = Cisco Systems (with Cisco Intrusion Prevention Security Agent)
Required-Client-Firewall-Product-Code	Y	Y	Y	46	Integer	Single	Cisco Systems Products: 1 = Cisco Intrusion Prevention Security Agent or Cisco Integrated Client (CIC) Zone Labs Products: 1 = Zone Alarm 2 = Zone AlarmPro 3 = Zone Labs Integrity NetworkICE Product: 1 = BlackIce Defender/Agent Sygate Products: 1 = Personal Firewall 2 = Personal Firewall Pro 3 = Security Agent
Required-Client-Firewall-Description	Y	Y	Y	47	String	Single	String
Require-HW-Client-Auth	Y	Y	Y	48	Boolean	Single	0 = Disabled 1 = Enabled
Required-Individual-User-Auth	Y	Y	Y	49	Integer	Single	0 = Disabled 1 = Enabled
Authenticated-User-Idle-Timeout	Y	Y	Y	50	Integer	Single	1-35791394 minutes

Table C-7 Security Appliance Supported RADIUS Attributes and Values (continued)

Attribute Name	VPN 3000	ASA	PIX	Attr. #	Syntax/ Type	Single or Multi-Valued	Description or Value
Cisco-IP-Phone-Bypass	Y	Y	Y	51	Integer	Single	0 = Disabled 1 = Enabled
IPSec-Split-Tunneling-Policy	Y	Y	Y	55	Integer	Single	0 = No split tunneling 1 = Split tunneling 2 = Local LAN permitted
IPSec-Required-Client-Firewall-Capability	Y	Y	Y	56	Integer	Single	0 = None 1 = Policy defined by remote FW Are-You-There (AYT) 2 = Policy pushed CPP 4 = Policy from server
IPSec-Client-Firewall-Filter-Name	Y			57	String	Single	Specifies the name of the filter to be pushed to the client as firewall policy
IPSec-Client-Firewall-Filter-Optional	Y	Y	Y	58	Integer	Single	0 = Required 1 = Optional
IPSec-Backup-Servers	Y	Y	Y	59	String	Single	1 = Use Client-Configured list 2 = Disable and clear client list 3 = Use Backup Server list
IPSec-Backup-Server-List	Y	Y	Y	60	String	Single	Server Addresses (space delimited)
DHCP-Network-Scope	Y	Y	Y	61	String	Single	IP Address
Intercept-DHCP-Configure-Msg	Y	Y	Y	62	Boolean	Single	0 = Disabled 1 = Enabled
MS-Client-Subnet-Mask	Y	Y	Y	63	Boolean	Single	An IP address
Allow-Network-Extension-Mode	Y	Y	Y	64	Boolean	Single	0 = Disabled 1 = Enabled
Authorization-Type	Y	Y	Y	65	Integer	Single	0 = None 1 = RADIUS 2 = LDAP
Authorization-Required	Y			66	Integer	Single	0 = No 1 = Yes
Authorization-DN-Field	Y	Y	Y	67	String	Single	Possible values: UID, OU, O, CN, L, SP, C, EA, T, N, GN, SN, I, GENQ, DNQ, SER, use-entire-name
IKE-KeepAlive-Confidence-Interval	Y	Y	Y	68	Integer	Single	10-300 seconds
WebVPN-Content-Filter-Parameters	Y	Y		69	Integer	Single	1 = Java ActiveX 2 = Java Script 4 = Image 8 = Cookies in images

Table C-7 Security Appliance Supported RADIUS Attributes and Values (continued)

Attribute Name	VPN 3000	ASA	PIX	Attr. #	Syntax/ Type	Single or Multi-Valued	Description or Value
WebVPN-URL-List		Y		71	String	Single	URL-List name
WebVPN-Port-Forward-List		Y		72	String	Single	Port-Forward list name
WebVPN-Access-List		Y		73	String	Single	Access-List name
Cisco-LEAP-Bypass	Y	Y	Y	75	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Homepage	Y	Y		76	String	Single	A URL such as http://example-portal.com
Client-Type-Version-Limiting	Y	Y	Y	77	String	Single	IPSec VPN version number string
WebVPN-Port-Forwarding-Name	Y	Y		79	String	Single	String name (example, “Corporate-Apps”). This text replaces the default string, “Application Access,” on the clientless portal home page.
IE-Proxy-Server	Y			80	String	Single	IP address
IE-Proxy-Server-Policy	Y			81	Integer	Single	1 = No Modify 2 = No Proxy 3 = Auto detect 4 = Use Concentrator Setting
IE-Proxy-Exception-List	Y			82	String	Single	newline (\n) separated list of DNS domains
IE-Proxy-Bypass-Local	Y			83	Integer	Single	0 = None 1 = Local
IKE-Keepalive-Retry-Interval	Y	Y	Y	84	Integer	Single	2 - 10 seconds
Tunnel-Group-Lock		Y	Y	85	String	Single	Name of the tunnel group or “none”
Access-List-Inbound		Y	Y	86	String	Single	Access list ID
Access-List-Outbound		Y	Y	87	String	Single	Access list ID
Perfect-Forward-Secrecy-Enable	Y	Y	Y	88	Boolean	Single	0 = No 1 = Yes
NAC-Enable	Y			89	Integer	Single	0 = No 1 = Yes
NAC-Status-Query-Timer	Y			90	Integer	Single	30 - 1800 seconds
NAC-Revalidation-Timer	Y			91	Integer	Single	300 - 86400 seconds
NAC-Default-ACL	Y			92	String		Access list
WebVPN-URL-Entry-Enable	Y	Y		93	Integer	Single	0 = Disabled 1 = Enabled

Table C-7 Security Appliance Supported RADIUS Attributes and Values (continued)

Attribute Name	VPN 3000	ASA	PIX	Attr. #	Syntax/ Type	Single or Multi-Valued	Description or Value
WebVPN-File-Access-Enable	Y	Y		94	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-File-Server-Entry-Enable	Y	Y		95	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-File-Server-Browsing-Enable	Y	Y		96	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Port-Forwarding-Enable	Y	Y		97	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Outlook-Exchange-Proxy-Enable	Y	Y		98	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Port-Forwarding-HTTP-Proxy	Y	Y		99	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Auto-Applet-Download-Enable	Y	Y		100	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Citrix-Metaframe-Enable	Y	Y		101	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Apply-ACL	Y	Y		102	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SSL-VPN-Client-Enable	Y	Y		103	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SSL-VPN-Client-Required	Y	Y		104	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SSL-VPN-Client-Keep-Installation	Y	Y		105	Integer	Single	0 = Disabled 1 = Enabled
SVC-Keepalive	Y	Y		107	Integer	Single	0 = Off 15 - 600 seconds
SVC-DPD-Interval-Client	Y	Y		108	Integer	Single	0 = Off 5 - 3600 seconds
SVC-DPD-Interval-Gateway	Y	Y		109	Integer	Single	0 = Off) 5 - 3600 seconds
SVC-Rekey-Time		Y		110	Integer	Single	0 = Disabled 1- 10080 minutes
WebVPN-Deny-Message		Y		116	String	Single	Valid string(up to 500 characters)
Extended-Authentication-On-Rekey		Y	Y	122	Integer	Single	0 = Disabled 1 = Enabled
SVC-DTLS		Y		123	Integer	Single	0 = False 1 = True

Table C-7 Security Appliance Supported RADIUS Attributes and Values (continued)

Attribute Name	VPN 3000	ASA	PIX	Attr. #	Syntax/ Type	Single or Multi-Valued	Description or Value
SVC-MTU		Y		125	Integer	Single	MTU value 256 - 1406 in bytes
SVC-Modules		Y		127	String	Single	String (name of a module)
SVC-Profiles		Y		128	String	Single	String (name of a profile)
SVC-Ask		Y		131	String	Single	0 = Disabled 1 = Enabled 3 = Enable default service 5 = Enable default clientless (2 and 4 not used)
SVC-Ask-Timeout		Y		132	Integer	Single	5 - 120 seconds
IE-Proxy-PAC-URL		Y		133	String	Single	PAC Address String
Strip-Realm	Y	Y	Y	135	Boolean	Single	0 = Disabled 1 = Enabled
Smart-Tunnel		Y		136	String	Single	Name of a Smart Tunnel
WebVPN-ActiveX-Relay		Y		137	Integer	Single	0 = Disabled Otherwise = Enabled
Smart-Tunnel-Auto		Y		138	Integer	Single	0 = Disabled 1 = Enabled 2 = AutoStart
Smart-Tunnel-Auto-Signon-Enable		Y		139	String	Single	Name of a Smart Tunnel Auto Signon list appended by the domain name
VLAN		Y		140	Integer	Single	0 - 4094
NAC-Settings		Y		141	String	Single	Name of NAC policy
Member-Of		Y	Y	145	String	Single	Comma delimited string, for example: `Engineering, Sales` This is an administrative attribute that can be used in dynamic access policies. It does not set a group policy.
Address-Pools		Y	Y	217	String	Single	Name of IP local pool
IPv6-Address-Pools		Y		218	String	Single	Name of IP local pool-IPv6
IPv6-VPN-Filter		Y		219	String	Single	ACL value
Privilege-Level		Y	Y	220	Integer	Single	An integer between 0 and 15.

Table C-7 Security Appliance Supported RADIUS Attributes and Values (continued)

Attribute Name	VPN 3000	ASA	PIX	Attr. #	Syntax/ Type	Single or Multi-Valued	Description or Value
WebVPN-Macro-Value1		Y		223	String	Single	Unbounded. See the *SSL VPN Deployment Guide* for examples and use cases at this URL: http://supportwiki.cisco.com/ViewWiki/index.php/Cisco_ASA_5500_SSL_VPN_Deployment_Guide%2C_Version_8.x
WebVPN-Macro-Value2		Y		224	String	Single	Unbounded. See the *SSL VPN Deployment Guide* for examples and use cases at this URL: http://supportwiki.cisco.com/ViewWiki/index.php/Cisco_ASA_5500_SSL_VPN_Deployment_Guide%2C_Version_8.x

Security Appliance IETF RADIUS Authorization Attributes

Table C-8 list all the possible IETF Radius attributes.

Table C-8 Security Appliance Supported IETF RADIUS Attributes and Values

Attribute Name	VPN 3000	ASA	PIX	Attr. #	Syntax/ Type	Single or Multi-Valued	Description or Value
IETF-Radius-Class	Y	Y	Y	25		Single	Sets the group policy for the remote access VPN session. For 8.2 and later, we recommend that you use the Group-Policy attribute. You can use one of the three following formats: • <group policy name> • OU=<group policy name> • OU=<group policy name>;
IETF-Radius-Filter-Id	Y	Y	Y	11	String	Single	Access list name that is defined on the security appliance. This applies only to full tunnel IPsec and SSL VPN clients
IETF-Radius-Framed-IP-Address	Y	Y	Y	n/a	String	Single	An IP address
IETF-Radius-Framed-IP-Netmask	Y	Y	Y	n/a	String	Single	An IP address mask

Table C-8 Security Appliance Supported IETF RADIUS Attributes and Values

IETF-Radius-Idle-Timeout	Y	Y	Y	28	Integer	Single	seconds
IETF-Radius-Service-Type	Y	Y	Y	6	Integer	Single	seconds. Possible Service Type values: .Administrative—user is allowed access to configure prompt. .NAS-Prompt—user is allowed access to exec prompt. .remote-access—user is allowed network access
IETF-Radius-Session-Timeout	Y	Y	Y	27	Integer	Single	seconds

Configuring an External TACACS+ Server

The security appliance provides support for TACACS+ attributes. TACACS+ separates the functions of authentication, authorization, and accounting. The protocol supports two types of attributes: mandatory and optional. Both the server and client must understand a mandatory attribute, and the mandatory attribute must be applied to the user. An optional attribute may or may not be understood or used.

Note To use TACACS+ attributes, make sure you have enabled AAA services on the NAS.

Table C-9 lists supported TACACS+ authorization response attributes for cut-through-proxy connections. Table C-10 lists supported TACACS+ accounting attributes.

Table C-9 Supported TACACS+ Authorization Response Attributes

Attribute	Description
acl	Identifies a locally configured access list to be applied to the connection.
idletime	Indicates the amount of inactivity in minutes that is allowed before the authenticated user session is terminated.
timeout	Specifies the absolute amount of time in minutes that authentication credentials remain active before the authenticated user session is terminated.

Table C-10 Supported TACACS+ Accounting Attributes

Attribute	Description
bytes_in	Specifies the number of input bytes transferred during this connection (stop records only).
bytes_out	Specifies the number of output bytes transferred during this connection (stop records only).
cmd	Defines the command executed (command accounting only).
disc-cause	Indicates the numeric code that identifies the reason for disconnecting (stop records only).

Table C-10 Supported TACACS+ Accounting Attributes (continued)

Attribute	Description
elapsed_time	Defines the elapsed time in seconds for the connection (stop records only).
foreign_ip	Specifies the IP address of the client for tunnel connections. Defines the address on the lowest security interface for cut-through-proxy connections.
local_ip	Specifies the IP address that the client connected to for tunnel connections. Defines the address on the highest security interface for cut-through-proxy connections.
NAS port	Contains a session ID for the connection.
packs_in	Specifies the number of input packets transferred during this connection.
packs_out	Specifies the number of output packets transferred during this connection.
priv-level	Set to the user's privilege level for command accounting requests or to 1 otherwise.
rem_iddr	Indicates the IP address of the client.
service	Specifies the service used. Always set to "shell" for command accounting only.
task_id	Specifies a unique task ID for the accounting transaction.
username	Indicates the name of the user.

INDEX